

REPORT DOCUMENTATION PAGE

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302, and to the Office of Management and Budget, Paperwork Project Director, Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE 26-Feb-99	3. REPORT TYPE AND DATES COVERED Demonstration Test Report
4. TITLE AND SUBTITLE Project Plan		5. FUNDING NUMBERS Contract N00421-97-C-1293/P4	
6. AUTHOR(S) Steven Case Paolo Dini			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) ViA Inc. 11 Bridge Square Northfield, MN 55057		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) & ADDRESS(ES) Charles D. Caposell Code 4.0T Naval Air Systems Command Building 2185, Ste 1190 22347 Cedar Point Rd. Unit #6 Patuxent River, MD 20670-1161		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION/AVAILABILITY STATEMENT		12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Project Plan #3			
14. SUBJECT TERMS		15. NUMBER OF PAGES	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	
20. LIMITATION OF ABSTRACT			

19990401 097

Standard Form (298)

DTIC QUALITY INSPECTED 2

AQU99-06-1218

Project Plan

Period Covered by the Report

1 November 1998 through 31 January 1999

Date of Report: 26 February 1999

Wrist Interactive Device for Wearable PC

SBIR Phase II Topic N95-137

Contract No. N00421-97-C-1293

Dollar Value \$1,708,653

ViA Inc.

11 Bridge Square

Northfield, MN 55057

Sponsor

Charles D. Caposell

Naval Air Systems Command

AIR-4.5T

Data Item No. 001

Contract Reference Item 0001

Authority - Data Acquisition Documentation No. DI-MGMT-80227

Project Plan Report No. 3

Issuing Government Activity

Requiring Office AIR-4.0T

Security Classification – Unclassified

DTIC STATEMENT

TITLE: Project PLAN

PLEASE CHECK THE APPROPRIATE BLOCK BELOW:

AO #499-06-1218

☐ 1 copies are being forwarded. Indicate whether Statement A, B, C, D, E, F, or X applies.

☒ DISTRIBUTION STATEMENT A:
APPROVED FOR PUBLIC RELEASE: DISTRIBUTION IS UNLIMITED

☐ DISTRIBUTION STATEMENT B:
DISTRIBUTION AUTHORIZED TO U.S. GOVERNMENT AGENCIES ONLY; (Indicate Reason and Date). OTHER REQUESTS FOR THIS DOCUMENT SHALL BE REFERRED TO (Indicate Controlling DoD Office).

☐ DISTRIBUTION STATEMENT C:
DISTRIBUTION AUTHORIZED TO U.S. GOVERNMENT AGENCIES AND THEIR CONTRACTORS; (Indicate Reason and Date). OTHER REQUESTS FOR THIS DOCUMENT SHALL BE REFERRED TO (Indicate Controlling DoD Office).

☐ DISTRIBUTION STATEMENT D:
DISTRIBUTION AUTHORIZED TO DoD AND U.S. DoD CONTRACTORS ONLY; (Indicate Reason and Date). OTHER REQUESTS SHALL BE REFERRED TO (Indicate Controlling DoD Office).

☐ DISTRIBUTION STATEMENT E:
DISTRIBUTION AUTHORIZED TO DoD COMPONENTS ONLY; (Indicate Reason and Date). OTHER REQUESTS SHALL BE REFERRED TO (Indicate Controlling DoD Office).

☐ DISTRIBUTION STATEMENT F:
FURTHER DISSEMINATION ONLY AS DIRECTED BY (Indicate Controlling DoD Office and Date) or HIGHER DoD AUTHORITY.

☐ DISTRIBUTION STATEMENT X:
DISTRIBUTION AUTHORIZED TO U.S. GOVERNMENT AGENCIES AND PRIVATE INDIVIDUALS OR ENTERPRISES ELIGIBLE TO OBTAIN EXPORT-CONTROLLED TECHNICAL DATA IN ACCORDANCE WITH DoD DIRECTIVE 5230.25, WITHHOLDING OF UNCLASSIFIED TECHNICAL DATA FROM PUBLIC DISCLOSURE, 6 Nov 1984 (Indicate date of determination). CONTROLLING DoD OFFICE IS (Indicate Controlling DoD Office).

☐ This document was previously forwarded to DTIC on _____ (date) and the AD number is _____.

☐ In accordance with provisions of DoD instructions, the document requested is not supplied because:

☐ It will be published at a later date. (Enter approximate date, if known).

☐ Other. (Give Reason)

DoD Directive 5230.24, "Distribution Statements on Technical Documents," 18 Mar 87, contains seven distribution statements, as described briefly above. Technical Documents must be assigned distribution statements.

STEVEN CASE

Print or Type Name

507 663 1399

Telephone Number

Joyce J. Chinos
Authorized Signature/Date

CDRL Distribution List and Addresses

Commander
Naval Air Systems Command
Attn: Mr. Charles Caposell, Code 4.0T
Building 2185, Suite 1190
22347 Cedar Point Rd., Unit #6
Patuxent River, MD 20670-1161

Department of the Navy
Naval Air Systems Command
Attn: Mr. Geoff Tisone, AIR-2.5.3
MS 34, Building 3189
Naval Air Station
Patuxent River, MD 20670-5449

Commander, Naval Air Systems Command
Attn: Technical Library, Code AIR-4.0C1
Building 2185
22347 Cedar Point Rd., Unit #6
Patuxent River, MD 20670-1161

Director
Defense Advanced Research Projects Agency / ETO
Attn: E.C. Urban
3701 North Fairfax Dr.
Arlington, VA 22203-1714

Defense Technical Information Center
Attn: DTIC-User Services
8725 Kingman Rd
Fort Belvoir, VA 22060-6218

Office of Special Technology
Attn: Kathleen Griggs
10530 Riverview Road
Fort Washington, MD 20744

CONTENTS

<u>Chapter</u>	<u>Page</u>
I. INTRODUCTION	1
II. SURVEY OF TECHNOLOGY	3
Display	3
a) Reflective/Transmissive	3
b) Electroluminescent	5
c) Field-Emission	7
RF Interface	10
Microprocessor	13
Microphone	16
Speaker	16
Audio CODEC	16
Pager/Vibrator	16
Power Budget Estimate	17
Voice/Image Compression and Data Reduction	17
Flexible Motherboard	18
Preliminary Block Diagrams of WID Components	19
III. DEVELOPMENT OF WID EVALUATION BOARD #1	22
Comparison of the Intel SA1100 and the NEC Vr4111 Microprocessors	22
Interfacing the SA1100 to the FED Display	27
Interfacing the SA1100 to the MicroDisplay Display	35
RF Interface	36
RAM and Flash	38
Board Schematics	38
IV. MAGNIFICATION OPTICS	53
Background and Preliminary Study	53
Design of Optics Mock-Up for Emissive Display	63
Design of Optics Mock-Up for Reflective Display	66
V. BATTERY TECHNOLOGY	72
Introduction	72
Lithium-Ion Batteries	74
Lithium-Polymer Batteries	75
Lithium-V ₂ O ₅ Batteries	75
Appendix: Review Of Fuel Cell Technology	81
VI. DISPLAY AND COMMUNICATION DRIVERS	83
Drivers on the Host Computer (ViA II)	83
High-Level CE Application on the WID	85
Porting Windows CE to the SA1100 Microprocessor	85

VII. DEVELOPMENT OF WID EVALUATION BOARD #2	89
Overview	89
Audio System Characteristics and Implementation Considerations	91
VIII. SCHEDULE	102
IX. PRELIMINARY PHYSICAL MODEL OF THE WID	103

Chapter I

INTRODUCTION

ViA, Inc., has conducted a search of existing technology for the development of a wrist interactive device (WID). This DARPA-funded device will provide a compact, wearable interface to ViA's wearable computers and other hardware platforms. ViA's wearable computers (such as the ViA II) are worn around the waist, whereas the WID will be worn on the wrist. The WID will have the following functionality:

- display of PC screen with VGA resolution or better
- two-way voice communication
- paging
- several hours of operation

To achieve this functionality, the WID will include the following components, also shown schematically in Fig. 1:

- miniature display with magnifying optics that provide a large virtual image
- speaker/microphone with corresponding audio CODEC
- RF interface chipset
- buzzer/vibrator
- map reader
- battery
- microprocessor to manage the entire system.

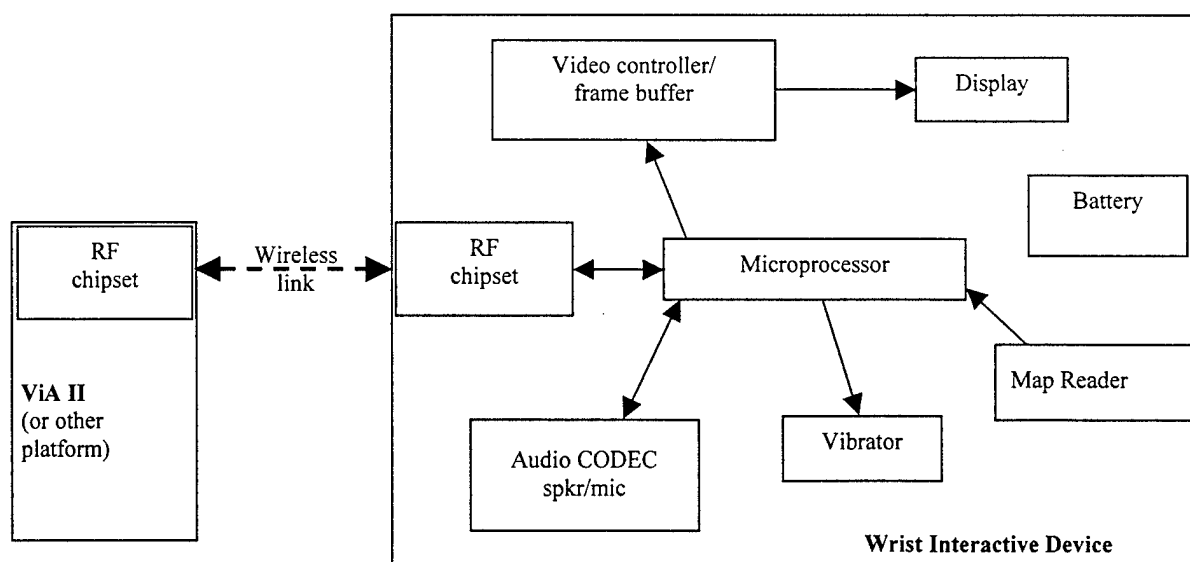


Fig. 1.1 Schematic of WID Components Showing Data Flow

Each of these components must be extremely small in size and have low power requirements. There is no combination of existing components today that satisfies these size, power, and performance requirements. However, ViA has identified a variety of products that come close and whose subsequent versions, under current development or in the next twelve months, will meet the specifications required by DARPA.

The rapid evolution of technology makes it possible to consider an alternative, more aggressive research and development path. IBM, National Semiconductor, and other companies are integrating an increasing number of components onto a single chip. Such system-on-a-chip solutions promise to incorporate an entire PC on a single chip in the near future, using flash memory for mass storage. The operating system for this computer will be Windows CE, which uses only 4 MB of RAM. The single-chip configuration would eliminate the requirement of an RF interface for the WID, in one possible version. The heart of the WID would be composed of several layers: a flexible lithium-polymer battery as the wristband; next, the BGA PC chip, "pins up," attached to the flex motherboard. On the other side of the motherboard will be an electroluminescent display, also made of several layers of plastic, with a flexible magnifying lens on top. Power requirements would be kept at approximately 10% of current levels by power management technology currently pursued by Intel, where the processor is turned off between instructions. As a consequence, overheating of the polymers would not be an issue and several hours of operation could be expected of the WID before recharging.

Over the next year ViA will keep examining both possible implementations of the WID. It is not expected that a full "wrist PC" will become possible before the expiration of the grant. Therefore, this report will discuss technology mostly relevant to the original WID configuration.

Chapter II discusses the available technology for each component. Chapter III details ViA's analysis and design efforts to develop the WID development board. Chapter IV gives an overview of the software drivers being developed to run on the StrongARM microprocessor. Roadmaps of the planned development stages are provided where applicable.

Chapter II

SURVEY OF TECHNOLOGY

Display

ViA has been involved with several display companies as part of its on-going product development and has contacted several companies that manufacture miniature displays that are potential candidates for inclusion into the WID. Prior to the start of the contract, ViA worked with Kopin, Displaytech, Planar, and Colorado MicroDisplay. Since the initiation of this contract, ViA has visited MicroDisplay's facilities (a different company from Colorado MicroDisplay). In May 1998, ViA also attended the Society for Information Display Conference in order to continue discussions with possible display vendors. Additional visits to these and other companies are planned. These displays vary in the thickness of the optics required to magnify and illuminate their image. There are also variations in the technology used for the pixels, and for handling or buffering the data to be displayed. The display technologies ViA has examined can be classified in terms of the illumination source: a) reflective/transmissive, b) electroluminescence, and c) field emission.

a) Reflective/Transmissive. MicroDisplay makes an extremely small, reflective, active-matrix LCD display that has VGA resolution and is capable of generating full color. One of the interesting features of the MicroDisplay device is the generation of color by the top layer of electrodes. These electrodes are configured as three different diffraction gratings, each regulated by applying voltages proportional to the desired RGB levels. The interface to the host computer translates an NTSC or VGA signal into the voltage levels to be applied to the gratings. This is masterfully done by a circuit integrated onto the display chip itself, so that the power and size savings are significant. In addition, the silicon-based display technology makes the manufacturing of this display very affordable. For incorporation in the wrist interactive device, however, this display has two drawbacks:

- 1) The present interfaces of the monochrome display are to an NTSC or a VGA analog video signal and have been integrated onto the display chip. Unfortunately, errors in the first design are preventing volume production of this display, for the moment. Since the host processor ViA is considering is capable of driving the display through a TFT interface, a RAMDAC chip could be used to convert the digital data to an analog RGB signal. However, the best solution for the WID would be a memory-mapped interface.
- 2) MicroDisplay has released a color display with a field-sequential color-generating method. The drive circuitry that this approach requires has been shrunk to an ASIC that, however, still uses 1 W of power, although the display chip itself is very low-power (15 mW). In addition, the frame buffer for MicroDisplay's ASIC is not integrated but consists of two external SRAM chips, which increase the space required by the electronics inside the WID. The color-generating method that relies on the diffraction gratings and that uses less power, although it has been demonstrated in the lab, has not been implemented yet and production units are not available.

Displaytech also has developed a reflective miniature display of 640 x 480 resolution. It uses a ferroelectric liquid crystal whose behavior in response to applied voltages is inherently digital, since it is either fully polarized or fully unpolarized. Red, green and blue LEDs are shined in sequence for the same time duration. Thus each pixel is either fully colored or black, while illuminated by each LED. To achieve different colors and a grayscale effect, the saturation of each 5-bit color for each pixel is translated into a pulse-width modulation of the voltage applied to the electrodes. In other words, the 5-bit binary number specifying the saturation of each color for each pixel determines the fraction of the time each pixel is "turned on" while being illuminated by the LED of that color. The overall effect of this sequential PWM mixing is indistinguishable from other forms of 15-bit color. The display requires the data in Low Voltage Differential Signaling (LVDS) format, which is a serial digital protocol carrying the PWM signal for each pixel. The translation from parallel digital to LVDS PWM is effected by a Field-Programmable Gate Array (FPGA) chip. At this point, such a chip requires a great deal of power and is rather large. Precise figures were not disclosed by Displaytech because they consider the FPGA chip only an evaluation board-level solution. OEMs such as ViA need to develop their own ASIC for their commercial applications, to perform the task of the FPGA. While an ASIC would clearly make possible significant savings in power and size, the cost of its development is very high. The recently announced alliance between Displaytech and Hewlett-Packard may lead HP to develop such an ASIC for this display. However, the ASIC would not be available until Q3 of 1999.

Colorado MicroDisplay (a different company from MicroDisplay) has developed a dynamic, nematic, liquid-crystal-on-silicon technology that takes advantage of very affordable and widespread CMOS and LCD manufacturing processes. The fast switching speed of the nematic liquid crystal allows up to 90 frames per second, where each frame is completed after flashing red, green, and blue LEDs in sequence. The switching speed of this crystal can therefore reach up to 270 Hz. During an initial visit by CMD at ViA in November 1998 and subsequent conference calls, it has become apparent that ViA and CMD share several common development goals. CMD is developing an ASIC to facilitate the task of interfacing their microdisplay to OEM host devices, and has expressed interest in supporting ViA's efforts in the WID project. Accordingly, CMD's microdisplay will present an interface to the WID's microprocessor that imposes the smallest load both in terms of MIPS as well as system bus bandwidth. Evaluation units of this display will become available in Q2 99. Although this display is reflective, by relying on the expertise of DisplayWear, an optical engineering consulting company, ViA is developing a collapsible magnification system that may satisfy the required form factor of the WID without sacrificing the functionality and readability of the CMD display. ViA and DisplayWear held a meeting with CMD at CMD's facilities in January 1999 to discuss the current status and the future development of the display system together with CMD's optical, mechanical, and system engineers.

Kopin has developed an active-matrix, 320 x 240 pixel, color, transmissive display. Here, the light comes from behind the display, as in a slide projector. Kopin has been able to fit the backlight module, the display, and the magnifying optics within an overall 2 cm thickness, which is better than most alternatives and may be acceptable for the WID. Kopin's interface is analog/digital and they achieve full color by refreshing the screen 180 times a second, 60 times

for each color ("field-sequential"). The analog color information for each pixel is translated into three voltage levels (8 bits each) that determine its degree of polarization while the three RGB LEDs illuminate it, in sequence. The digital lines carry clock and sync signals. Although the power consumption of the display apparatus is only 46 mW, the need for a frame buffer and a video controller is problematic for the WID's power budget. Two things make Kopin's display still interesting to ViA: 1) Kopin is planning to release a 640 x 480 display in 1999, and 2) the eye relief of the Kopin display is significantly longer than that of most other reflective displays (3 inches). ViA will therefore keep investigating Kopin's display as a possible alternative solution.

Three-Five has developed an SVGA (800 x 600), monochrome, active-matrix, LCD-on-silicon display technology that, like most of these processes, allows pixels smaller than what the human eye can resolve. Also their screen is reflective and therefore requires front illumination in addition to magnification. Three-Five has licensed their display to Siliscape, who has developed a compact magnification and illumination optical system for it, leading to the Optiscape display. Optiscape's packaging dimensions are only 30 mm x 36 mm with an overall thickness of 17.5 mm, including magnification and illumination. The optics of the Siliscape system are discussed in detail in Chapter IV of this report. The interface to the host of the Three-Five display is more desirable than other displays since it is digital and each pixel is composed of SRAM that only needs to be refreshed approximately once a second, and the microprocessor can write directly to it. The power draw of the display itself (100 mW) is higher than MicroDisplay's, but the absence of any drive circuitry associated with field-sequential color, a frame buffer, or an LCD controller makes the overall power requirement by far the lowest of all other alternatives. The Three-Five display, therefore, matches the WID requirements. Further assessment of the compatibility of this display to the WID system, as well as information on when a color version may become available, has been difficult to obtain since up to now Three-Five has not expressed any interest in working with ViA.

b) Electroluminescent. Another promising approach is based on electroluminescence. Here, two layers of transparent semiconducting material are doped in a manner analogous to a diode. A voltage is applied that causes the recombination of the free electrons with the holes, at the junction of the two layers, and energy is released in the form of light. ViA has looked at three companies developing this technology: Cambridge Display Technology (CDT), Planar Systems, and FED Corporation. The first has pioneered the use of flexible polymers, which is of greatest interest to ViA, the second produces rigid screen panels, and FED has developed an organic light-emitting diode technology.

The technology currently available from CDT is analogous to passive matrix displays (although active matrix implementations are being developed), with two mutually perpendicular arrays of transparent electrodes mounted on either side of the semi-conducting polymer sheets. The substrate can be made of flexible plastic. As two electrodes become active, during the screen scan, only the overlap region will become conducting and emit light. One of the most interesting features of this approach is that by appropriate doping it is possible to produce semi-conducting polymers of different band-gap energies. Upon recombination of the charges, therefore, light of different wavelengths is emitted by the different polymers, which is perceived as different colors. CDT has incorporated these different polymers in the same screen as different pixels, thereby

achieving full RGB color. Since the polymer material acts as an insulator when not subjected to a voltage, these pixels can be small and close to one another without any danger of leakage between them; thus, very high resolutions can be achieved. Because the light comes from within the screen, a miniature display and magnifying optics based on this technology can be made very thin. The power consumption of CDT's displays is approximately 10 mA/cm^2 . Assuming a $2.5 \text{ cm} \times 3 \text{ cm}$ screen, that gives 75 mA. To achieve a luminance bright enough for outdoors use requires the display to be driven at 5 V, yielding 375 mW. Unfortunately, the circuitry to interface such a miniature display to a microprocessor has not been developed yet and would have to be pursued as a collaboration between ViA and CDT. Although it is difficult to estimate the power requirements of such interface circuit, 200 mW is not unreasonable if a custom, digital solution could be developed. Thus, we get approximately 600 mW for the overall display system. CDT has emphasized that the efficiency of their display technology is improving, so that this could be a conservative figure.

Philips has recently announced that they have licensed CDT's technology and will also manufacture this type of LEP displays, but their pilot project will deliver the first products some time in 1999 and it will then be only a 100×80 resolution display. Although CDT has expressed interest in working with ViA to develop an LEP display to ViA's specifications, both this and the Philips option require significant resources and development time. Therefore, LEP displays will be taken into consideration only if an easier solution cannot be made to work.

Secondly, Planar Systems, Inc., has developed a miniature, silicon-based, electroluminescent display, the MicroBrite. It is very small and thin, has full color and VGA resolution, and is active-matrix. Unfortunately, it uses 2 W of power and requires a VGA controller with frame buffer, further increasing the power requirements (by 1 W, approximately), and does not have a magnification system developed. The reason for the much higher power consumption of this type of display appears to originate from the semi-conducting materials. Unlike for CDT's polymer sheets, the SrS:Ce/ZnS:Mn semiconductors are prone to leakage among adjacent pixels when a voltage is applied between the electrodes. This is partly a consequence of the high voltages required to achieve electroluminescence with these materials (110 V for blue and 190 V for white light) and is remedied by inserting two thin insulating sheets between the phosphorus and the electrodes. The insulating property of the sheets breaks down temporarily and locally when the high voltage is applied. Since it is more difficult than with polymers to separate different colors by different amounts of doping, Planar relies on a subtractive, sequential, liquid-crystal color shutter (essentially a polarizing filter) to shut out different colors from the emitted white light in rapid succession. Because this process is subtractive, only 10% of the light generated by the phosphor reaches the viewer—hence the high energy consumption of this display. Another concern associated with the high power consumption is the heat generated by this display, which reaches 100°C during operation. Although the display is very thin and this heat can therefore be conducted away fairly easily, using this display the WID would most probably operate at a temperature that could be uncomfortable for the user.

Thirdly, ViA has contacted FED Corporation. In spite of the name of this company, which is an acronym for Field-Emission Display, the miniature display they are developing is an organic LED and, therefore, an electroluminescent display. Of all the emissive display technologies, this

display appears to be the most compatible with the WID requirements and schedule. Since it does have some drawbacks, a list of positive and negative features follows:

Positive Features:

- emissive: requires no independent illumination source
- current monochrome model has a parallel digital interface (8 bits/pixel, 256 gray levels)
- relatively large (0.75" diagonal): magnifying optics correspondingly thinner
- high resolution (1240 x 1024); VGA resolution easily programmable
- active-matrix
- collapsible prototype of magnifying optics developed by DisplayWear for ViA

Negative Features:

- high power draw (1.25 W)
- requires two voltages, 5 (digital circuit) and -7 to 12 V (cathode). The power supply circuit would therefore be complex and more likely to be inefficient.
- currently monochrome, color available by 2000
- color model is likely to have an analog interface
- monochrome display is in development and not yet available

While in principle CDT's LEP display is potentially better than FED's, the advantage of the latter is that it is near the end of its development and is scheduled to become available for evaluation later this year, complete with its digital interface and supporting electronics.

ViA will continue monitoring the development of all electroluminescent display technologies because of their thin form factor, but will focus on the FED display as the most promising of their kind.

c) Field-Emission. Finally, the last technology ViA has examined are field-emission displays. Conceptually, a field-emission display is equivalent to a CRT display where, however, the distance between the emitter and the phosphor has been shrunk to a fraction of a millimeter and where each pixel is treated as a separate cathode-anode pair.

Micron Technology makes a 548 x 222 pixels monochrome display that uses 350 mW. Its interface is to an NTSC signal, which is incompatible with a microprocessor. This display is not considered compatible with the WID.

Table 2.1 summarizes the main features of the displays that ViA has examined.

Table 2.1 MINIATURE DISPLAYS

	Resolution	Power ¹ (mW)	V	Size (mm)	Interface	Color	Comments
Reflective							
Displaytech	640 x 480	130+1000	5	8.32 x 6.24	VGA (Analog/Digital ²)	15-bit RGB Sequential	Needs front light source Drive ASIC/frame buffer not offered.
MicroDisplay	640 x 480	80 + 200	5	10.1 x 7.5	VGA (Analog/Digital ²)	Monochrome grayscale	Display not currently available. Field-sequential color display provided with ASIC that uses 1 W.
Three-Five (Siliscap)	800 x 600	200	5	?	Memory-mapped	3-bit mono grayscale	No frame buffer or video controller needed.
Colorado MicroDisplay	800 x 600	300 + 600	5.4	9.6 x 7.2	Digital TFT and Memory-mapped	18-bit RGB Sequential	Low-cost CMOS & LCD process. Drive ASIC/Frame buffer provided.
Transmissive							
Kopin	320 x 240	80 + 1000	3.3	4.8 x 3.6	VGA (Analog/Digital ²)	24-bit RGB Sequential	Needs back light source Needs frame buffer if processor is not fast enough to paint display
Electroluminescent							
Cambridge Display Technology	Any resolution	375+1000	3-5	tbd	Tbd	Monochrome	Needs no independent light source Light-Emitting Polymer display May need frame buffer
Philips	100 x 80	Tbd	3-5	Tbd	Tbd	Monochrome	LEP display, under development
Planar Systems	640 x 480	2000+1000	7	15.5 x 11.4 (0.75" diag.)	VGA (Analog/Digital ²)	RGB Subtractive	Active matrix EL, requires frame buffer
FED Corporation	640 x 480	1250	1 2	15.4 x 12.3 (0.77" diag.)	Digital TFT 35 // lines	Mono (now) RGB (1999)	No frame buffer beyond that used by the SA1100's LCD controller. Display not available yet.
Field-Emission							
Micron Technology	548 x 222	350+1000	5	0.55" diag.	NTSC	Monochrome	Needs no independent light source No plans to develop VGA/digital display

1 = The first number is the display itself + the illumination. The second number, if present, is the power dissipated by the drive circuitry, ASIC, frame buffer, RAMDAC, and/or PLL. A nominal 1000 mW is assumed for the drive ASIC unless more specific information has been provided. Finally, notice that a memory-mapped interface brings a system power savings of approximately 100 mW, since the SA1100's LCD controller would not need to be used.

2 = "Analog" refers to the data, i.e. the voltages required to achieve the desired gray levels, "Digital" refers to the sync and clock signals

Table 2.1 MINIATURE DISPLAYS (Continued)

	Medium	Brightness (cd/m ²)	Form Factor with ViA Optics	Interface ASIC?	Heat Dissipated	Operating Temperature (°C)	Frame Rate (Hz)	Available Now: Display/ASIC
Reflective								
Displaytech	Ferroelectric LC on CMOS	680	OK	NO	Low	0 to 45	60	Yes/No
MicroDisplay	Nematic LC on CMOS	?	OK	N/A	Low	0 to 45	60	No/ N/A
Three-Five (Siliscap)	SRAM	?	OK	N/A	Very Low	-10 to 60	N/A	Yes/ N/A
Colorado MicroDisplay	Nematic LC on CMOS	100	OK	YES	Low	0 to 50	85	Yes/Yes
Transmissive								
Kopin	Twisted Nematic LC	75	OK	NO	Low	0 to 60	60	Yes/No
Electroluminescent								
Cambridge Display Technology	Light-Emitting Polymer (LEP)	?	Good	Tbd	?	?	?	No/No
Philips	LEP	?	N/A	?	?	?	?	No/No
Planar Systems	SrS:Ce/ZnS:Mn	34	Good	NO	Very High	-20 to 50	60	Yes/No
FED Corporation	Organic LED	1000	Good	N/A	High	-30 to 70	120	No/ N/A
Field-Emission								
Micron Technology	Mini-CRT	50	Good	NO	?	-30 to 80	Full video	Yes/No

RF Interface

Several companies were contacted regarding the RF interface. The starting point was an IEEE study group attended by several companies involved in wireless communications and the development of Wireless Personal Area Networks (WPANs). The purpose of the study group is to draft a proposal for a new IEEE standard that addresses the requirements of a WPAN. WPANs currently fall under the guidelines of, but are incompatible with, the IEEE 802.11 Wireless Local Area Network (WLAN) standard. IEEE 802.11 is a specification of the medium access control (MAC) and physical layer (PHY) of radios operating in the 400 MHz to 2.4 GHz band under 10 W of transmit (radiated) power. These specifications are appropriate for wireless networks of several hundreds of feet in range, with coexistence capabilities with other WLANs, and with data rates of 1-10 Mbps. Unfortunately, these physical performance characteristics, coupled with the complexity of the logic associated with spread-spectrum encoding, lead to high power consumption and large form factors. The WPAN study group was constituted to develop specifications for a 1-10 meter range radio of 1 mW radiated power, with moderate interference resistance capability, few channels, and frequency agility for collision avoidance. In other words, if one channel experiences interference, a low-power radio would dynamically switch to a different channel as needed. This type of solution has modest power requirements and can be manufactured to fit on a very small board. The main drawback of this approach is smaller data rates, on the order of 100 kbps.

Xetron Corporation has already started the development of such a low-power, frequency-agile radio, the Hornet transceiver. Xetron will complete by April 1999 an RF board comprised of the RF and A/D chip, a digital ASIC, and a 4 MHz microcontroller, all on 1" x 1.5" board. ViA may not need to use this controller, as there will be one already on the WID, and Xetron is planning to integrate their chipset further. This indicates that by 4Q of 1999 a small enough radio should be available, hopefully a single-chip solution. The current data rate of the Xetron radio is 38 kbps. Once the on-board controller is bypassed and the WID's RISC processor is used instead, the data rate will increase to 115 kbps. A further increase can be achieved by taking advantage of the smaller 1-m range required by the WID relative to the current Hornet specification of 10 m. The bit duration can in fact be decreased and it may be possible to decrease the 5 samples/bit currently used by the Hornet to 3 samples/bit. The bit-rate should thus increase to, approximately, 400 kbps with the same probability of bit error. Its power draw is 130 mW in full receive/transmit mode and 33 mW daily average assuming 10% utilization.

Another, even better possibility for a short-range, low-power radio is the recently announced Bluetooth technology. Intel, Nokia, Toshiba, IBM, and Ericsson formed a special interest group and agreed on the performance specifications for a new radio to connect up to 8 computers (masters) or up to 255 active peripherals (slaves) within a 10 m range. Bluetooth is expected to require only 100 mW per device to operate and to achieve a data rate of 720 kbps. Additionally, using a time-domain multiple-access scheme (TDMA), the Bluetooth specification also supports 3 voice channels per frame for direct voice-to-voice communication between 2 active devices; i.e., cell phone, voice over IP, remote speakers/microphones, etc. Its form factor is fitting for the WID, as the entire radio will be integrated onto a 9 mm x 9 mm x 1.5 mm chip. Its only drawback is that production units won't be available until the end of 1999. ViA is slated to

receive Bluetooth development boards as soon as they become available, which is expected to be sometime in Q2 '99.

A possible alternative to a Bluetooth radio is being developed by Aeptec Microsystems, Inc. Aeptec is a sensor and sensor interfaces manufacturer and has developed an impressive single-chip transmitter. This radio broadcasts on the 900 MHz band with a direct-sequencing spread-spectrum MAC, using only 80 mW and achieving a raw data rate of 2 Mbps. The actual system data rate may be estimated around 1 Mbps but will depend on the implementation. It can coexist with up to 100 other similar radios. Although it is more complex than a frequency-agile system, this radio does not incur the overhead of the 802.11 Standard. The chip that achieves all this is essentially an ASIC that combines the digital and analog processing of the data stream in a tiny 3 mm-square package. Required supporting circuit components such as filters will increase the size of the final board. The current temperature range of the Aeptec transmitter is adequate for the WID prototype: 0-50 °C. Unfortunately, Aeptec has not developed the receiver, yet, although they estimate a 6-month delivery cycle. If Aeptec completes the development of their single-chip transceiver in a timely fashion, it will clearly be the best solution for the WID among all the options that ViA has examined.

A different approach altogether for wireless connectivity is to use medium-speed (MIR) and fast-speed (FIR) infrared technology. Such transceivers have variable data rates, from 115 kbps up to 4 Mbps, and low power consumption but are somewhat limited by line-of-sight requirements. This tends not to be a problem indoors, where the signals can bounce off walls and ceiling; outdoors, however, only the ground provides a reflective surface. Newer, 16 Mbps, wide-angle versions are currently under development at Hewlett-Packard, Siemens Microelectronics, and IBM Microelectronics Center. Engineering samples of these devices will be released after August 1999. While the wide-angle feature will improve the connectivity, outdoor use remains problematic also due to sunlight that can saturate the IR detector and pose a serious interference problem. Various narrow-band IR filters are readily available, but such polished optics are too expensive for implementation in most applications, for now. Finally, since the IrDA 1.0 and IrDA 1.1 specifications do not directly address a "multiple access" scheme, neither allows for multiple peripherals and computers to operate within the same space simultaneously. As a consequence, IR links are not considered optimal for the development of the WID and will not be investigated as thoroughly as RF links.

The characteristics of radios planned or already developed by these and other companies are listed in Table 2.2.

Table 2.2 RF AND IR INTERFACES

Manufacturer	Range (m)	Data Rate (kbps)	Power Draw max/avg (mW)	Trans. Power (mW)	Volt. (V)	Size	Carrier (GHz)	No. of Chann	Bandwidth/ Channel (MHz)	Medium Access Method	Modulation	Spreading
Xetron	10	115	130/33	1	3.3	1" x 1.5"	2.4	4	1	CSMA	Manchester OOK	Agile
AMP M/A Com	15	40	(not built yet)	0.75	-	-	2.4	556	0.150	-	Manchester FSK	FH 50 hops/s
GTE/BBN	2	130	20	1	3.3	-	0.900	8	1.5	TDMA	FSK	FH
Proxim	150	1600	1500/750	10	5	1.6"x2.6"	2.4	15	-	CSMA/CA		FHSS
Symbol	15	50	10	1	3.3	1" x 3"	2.4	79	1	CSMA/CA	FSK	Agile
Intermec Technologies	10	100	20	1	3.3	1.5 in ³ (PC card)	2.4	-	-	CSMA/CA	-	FHSS
Harris Semiconductor	100	1000	2500/1500	100	3.3	2.1" x 3.3"	2.4	-	-	CSMA/CA	Selectable D/BPSK D/QPSK	DSSS
HP IrDA	1	4000	15	25/Sr	3.3	4x4x10mm	500 kHz 1.152 MHz	N/A	N/A	ynchronous	Quad State PPM Sharp ASK	N/A
Aeptec (transmitter only, so far)	6	1000	80	2	3.3	1" x 1"	0.900	N/A	N/A	-	FSK	DSSS
Bluetooth	10	721	100	1	3.3	9x9x1.5mm	2.4	79	1.0	TDMA	2 & 4-level FSK	FH: 1600 hops/s

Microprocessor

The number of microprocessors available for controlling the WID is quite large, compared to the other components. It is helpful to classify the different types of microprocessors in terms of their computational power. In fact, there is a loose inverse relationship between the data rate of the RF link to the host computer and the computing power of the microprocessor on the WID. Two possible extremes are an RF link of unlimited bandwidth, requiring simply a multiplex/demultiplex function of the microprocessor; and a very modest RF data rate sending over, for instance, Postscript commands and requiring a sophisticated 32-bit processor for their interpretation. Although ViA will examine a number of configurations between these extremes, the current scenario leans toward higher processing power. Fig. 2.1 shows a possible breakdown of the processing tasks of the WID.

Of all the WID components, the microprocessor promises to be the least problematic. In fact, the size, power requirements, and heat dissipation of microprocessors are decreasing very rapidly, while their computational power is increasing just as rapidly. The huge and growing market of mobile computing and cellular phones has motivated large semiconductor houses like Motorola, Intel, National Semiconductor, and Hitachi to push this direction of development aggressively. As a result, it is feasible that the WID will eventually utilize a 32-bit RISC processor of at least 200 MHz clock speed and, at most, 200 mW power draw with internal ROM and RAM. Such a processor will be able to perform DSP functions for the speaker and microphones, compress/decompress the voice and display data, and (if necessary) act as the video controller for the display.

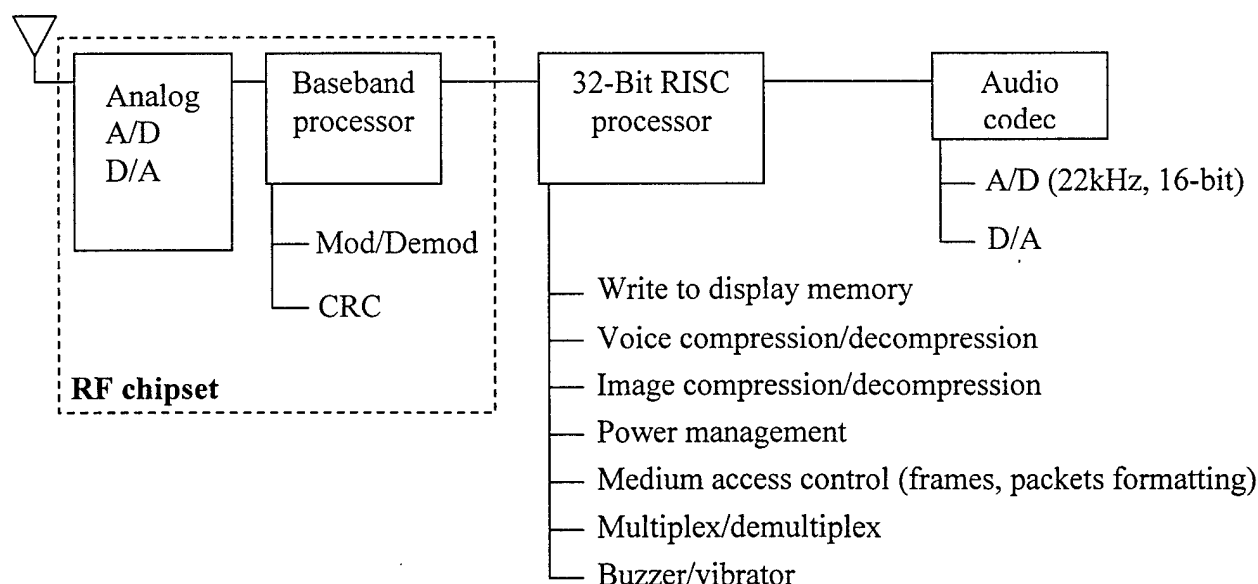


Fig. 2.1 Possible Breakdown of WID Processing Tasks

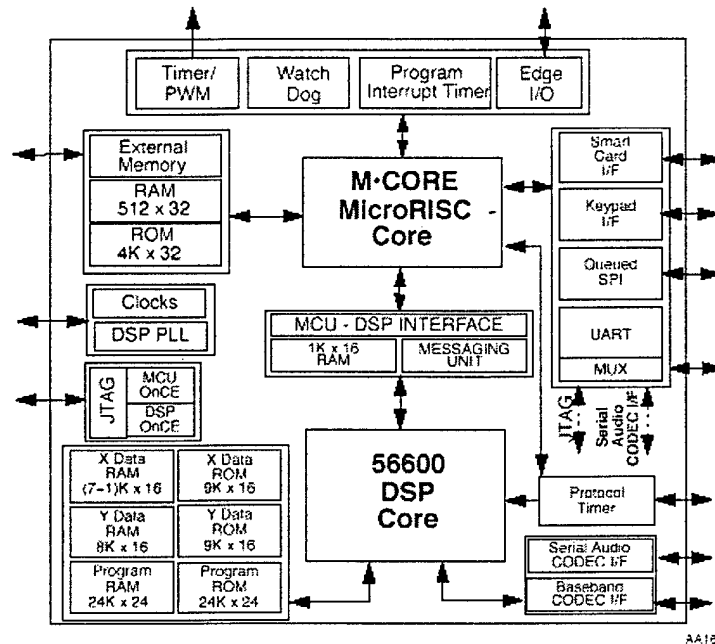


Fig. 2.2 Block Diagram of Motorola DSP56652 Processor.

Table 2.3 summarizes the principal brands and models that are being investigated. Among all these, three 32-bit processors seem better suited for the WID: the Intel SA1100 StrongARM processor, the Motorola DSP56652, and the NEC Vr4111. The SA1100 has an impressive performance, 230 MIPS for 300 mW, and an integrated LCD controller. The DSP56652 combines two very desirable features: a DSP and a 32-bit processor in one chip; moreover, the processor is based on the new Motorola M-CORE architecture that was developed with built-in power management specifically for mobile platforms. Fig. 2.2 shows a block diagram of this processor. Finally, the Vr4111 delivers a respectable 130 MIPS for 250 mW. However, the NEC processor necessitates a companion LCD controller chip that uses 225 mW, so the MIPS/Watt ratio is more favorable for the SA1100.

Table 2.3 MICROPROCESSORS

Manufacturer	Model	Family	Freq. (MHz)	MIPS	Process (microns)	Voltage (V)	Power (mW)	Package	Peripheral
AMD	ELANSC400	X86	100			2.7int/3.3ext	2046typ	292 BGA	
ARM	720T								
Cyrix	Gxi	X86	180	?	0.5u	2.9int/3.3ext	5550typ	352 BGA	5520
Cyrix	Gxm	X86	266	?	0.35u	2.9int/3.3ext	5870typ	352 BGA	5530
Fujitsu	MB86833	Sparclite	66	66	0.35u	3.3int/5ext			
Hitachi	SH7707	SH3	60	60	0.5u	3.3	560 typ	144 LQFP	"SOC"
Hitachi	SH7708R	SH3	100	100	0.5u	3.3	480 typ	144 QFP	
Hitachi	SH7709	SH3	100	100	0.5u	3.3	330 typ	208LQFP	HD64461
Hitachi	SH7750	SH4	200	360	0.35u	2.5 core	? typ	256LBGA	
IDT	WinChip	X86							
Intel	Pentium II Mobile	X86	233		0.25u	1.7 core	12100 max		
Intel	SA110	ARM	160	185	0.35u	1.65 core	450 typ	144LQFP	
Intel	SA1100 ¹	ARM	200	230	0.18u?	1.5 core	250 typ	208TQFP ²	
Mips	see NEC, Philips, Toshiba								
Motorola	DSP5652 ¹	Mcore/DSP	40/70	70	0.35	1.8/2.5	150	196PBGA 15x15	
Motorola	MMC2001	Mcore	33	31	0.36u	1.8int/3.3ext	80int max	144LQFP	
Motorola	MC68328	Dragonball							
Motorola	MPC821	PowerPC	50	66		2.2int/3.3ext	200 typ-2.2V	357 BGA	
Motorola	MPC823	PowerPC	50	66		2.2int/3.3ext	200 typ-2.2V	256PBGA 23x23	
NSC	DECT	Compact RISC							
NEC	Vr4310	Mips ISA III	167	221	0.28u	3.3	1500 typ	120PQFP 28x28	
NEC	Vr4111 ¹	" ISA III & M16	100	130	0.25u	2.5 core	200 typ	224FPBGA	
Philips	PR31700	Mips ISA II	75	72	?	3.3	110 typ	208PQFP 28x28	
Toshiba	TMPR3912U	Mips ISA II	75	75	0.35u	3.3	330 typ	208LQFP	
Toshiba	TMPR3922U	Mips ISA II	166	?	0.25u	2.5 core	400 typ	208LQFP	TC35143F

¹ = these processors appear to be close to the WID requirements

² = the SA1100 will be available in a BGA package, 17 mm x 17 mm, by the end of 1998

Microphone

Andrea Electronics Corporation has just developed a miniature noise-canceling microphone, 6 mm in diameter and 2 mm thick, that would fit quite easily on the wristband or on the side of the WID housing. Alternatively, a triangular array of three microphones could be used to filter out all out-of-phase sounds, so that only sounds originating along a line perpendicular to the plane of the WID would be passed. This is the most natural position for the user when issuing a voice command to the WID.

A possibly better approach is to utilize a tiny, 3 mm square microphone marketed by Noise Cancellation Technologies and manufactured by Siemens in conjunction with NCT's noise, feedback, and echo-canceling software. Feedback and echo canceling is important given the close proximity of the microphone to the speaker. The Siemens microphone will not be available until Q2 of 1999 but in the meantime ViA could start porting NCT's software to the selected microprocessor platform. NCT's software uses roughly 5 MIPS, depending on the platform it is ported to.

Speaker

To date, the smallest speaker ViA has identified is a mylar speaker with a plastic frame made by Mouser. With a diameter of 0.8" x 0.13" thickness, it will just fit in the wristband. This 8 Ω speaker draws 100 mW of power and has a frequency response of 550 Hz to 7 kHz, so it is appropriate for voice reproduction. Alternatively, an ear-phone bud may be connected to the WID by a retractable wire. By holding this speaker bud between the fingers of the hand wearing the WID, the hand would be placed over the ear and the user could then talk into the WID's microphone on the wrist, in front of his/her mouth. This will result in the same functionality as a telephone.

Audio CODEC

In order to achieve the resolution required for voice recognition, it is necessary to utilize a high-performance A/D converter and therefore a high performance audio CODEC such as the Texas Instruments TLC 320AD50. This CODEC samples the analog signals at 22 kHz with 16 bits of vertical resolution, has a programmable input/output gain control (important when background noise levels change), and uses 120 mW at full utilization. Although these characteristics match the WID requirements, this CODEC requires a 5 V power supply. A good alternative is the Analog Devices AD73311 which achieves the same resolution and programmability options with a 3 V supply.

Pager/Vibrator

Orion Devices Ltd., of England markets a possible vibrator motor for silent paging. Its size is 5 mm in diameter by 23 mm long (3/16" x 7/8"), with an operating voltage of 1.3 V, and current consumption of 80 mA. The power consumption is therefore about 100 mW, but the power draw averaged over one hour or longer is essentially zero.

Power Budget Estimate

Although the power budget of the WID is 500 mW, currently available components cannot fit within this figure. Thus, having examined most of the various components of the WID, the current power budget is estimated as follows ("average" assumes a 50% duty cycle for some components).

Table 2.4 Power Budget Estimate		
Component	Average Power Consumption (mW)	Maximum Power Consumption (mW)
RF Interface (Xetron)	75	130
• FED B&W display + CPLD	700	1400
• CMD color display + ASIC + LEDs	(450)	(900)
• MicroDisplay B&W display ¹ + LED + integrated VGA decoder + PLL + RAMDAC	[150] {100}	[300] {200}
• Three-Five B&W display + LED		
Microprocessor (SA1100)	300	800
Audio CODEC (AD73322)	60	120
DSP (ADSP2185L)	75	150
Speaker (Mouser)	50	100
Vibrator (Orion)	0	100
Map Reader		
Total Projected Power Consumption:		
with FED display	1260	2800
with CMD display	(1010)	(2300)
with MicroDisplay display	[710]	[1700]
with Three-Five display	{660}	{1600}

¹ = Currently not in production

Voice/Image Compression and Data Reduction

Downloading display information for a SVGA or even just a VGA display takes a very large bandwidth. For instance, to refresh a VGA screen in full RGB (24-bit) color 30 times a second requires 221 Mbps. This is clearly well beyond the capabilities of an RF or IR link. Lowering the refresh rate to once per second still requires 7.3 Mbps. Hewlett-Packard, at 16 Mbps, has demonstrated the fastest prototype IR link, while the fastest low-power radio promised by the Bluetooth consortium will deliver 720 kbps. Therefore, it would appear that an IR link is the obvious choice to achieve the required data rate. However, due to the absence of reflective surfaces outdoors, IR links are not as reliable as RF links. In either case, the use of data compression will be necessary to minimize bandwidth requirements and thus maximize overall system performance.

ViA has contacted Voxware, Inc., a leading firm developing voice and audio compression software. Voxware has experience developing low-level software to run on DSPs such as the TI C54x or the Motorola 5630x, and on RISC microprocessors such as the Hitachi SH-3 or the AMD ARM7. Because of the nature of ViA's wearable computers, most of its software is voice-enabled and relies on voice recognition engines such as Lernout & Hauspie's. Therefore, ViA implements high-resolution A/D conversion at 22 kHz sampling and 16-bit vertical resolution. Such resolution is sufficient to guarantee a high recognition rate on the part of the L&H ASR 1600 engine after the voice has been compressed, transmitted from the WID to the belt computer by the RF link, and decompressed. Voxware has codecs (compression-decompression software) that can deliver such voice data using only 6 kbps of bandwidth and a sizable 13 Mips of the WID processor.

As an alternative to Voxware, L&H itself develops also voice compression software and is interested in working with ViA. Although the bandwidth of L&H's compressed data stream is higher than Voxware's, there are significant advantages for ViA to working with L&H deriving from a recently forged partnership between these companies. For instance, the technical support in porting their software to ViA's hardware platforms generates real savings in human resources that ViA would otherwise have to employ for the same task.

Summus Technologies, Inc., is a good candidate for image and video compression. Their software is based on wavelet mathematics and can achieve 2 to 3 times more efficient compression than JPEG. While their compression algorithms will be indispensable for future versions of the WID for real-time video, it is not clear at this point whether image compression will be required for a standard graphical interface to a Windows application.

An alternative to data compression is data reduction. In this strategy, graphical commands are broadcast instead of low-level graphical information. For instance, if a window needs to be moved on the screen, the new coordinates could be sent to the WID instead of a new screen dump of all the screen pixels. In essence, this implies moving some of the device driver functions from the belt PC to the WID RISC processor. The format of these commands will depend on how much processing the WID RISC processor can accomplish. The overriding goal is to minimize the amount of data to be transferred over the RF link, since it has limited bandwidth. The specific implementation will be determined by how many MIPs and how much memory remain to the RISC processor to interpret display commands, after accounting for the RF baseband processing and the voice data compression and decompression.

Flexible Motherboard

It will be advantageous to rely on a flexible motherboard for the implementation of the WID circuitry. This is especially true since it appears that some of the components will be mounted on the wristband.

ViA has experience using polyamide-based flex boards. Its first and second-generation wearable computers use flex to connect rigid boards, while the third-generation computer will be entirely

made on flex. ViA has contacted 3M, who has just finished developing a new process for manufacturing flex circuits. Although their product is not commercially available in large quantities, yet, they can deliver up to 10 signal layers at 35μ /layer in small quantities. Their flex board has an impedance of $50\ \Omega$, $150\text{-}\mu$ capture pads, $300\text{-}\mu$ pad-to-pad spacing, $40\text{-}\mu$ lines, and $40\text{-}\mu$ spaces between lines. It would therefore satisfy the complexity requirements of the WID ICs in the most convenient package.

Preliminary Block Diagrams of WID Components

The search for components suitable for the WID is narrowing down. The principal criterion followed in the selection process is whether the part under examination has a projected performance by Q3 1999 that meets the WID requirements. Figs. 2.4 and 2.5 show block diagrams of two possible WID architectures using current choices for the components. Where appropriate, two alternatives are shown.

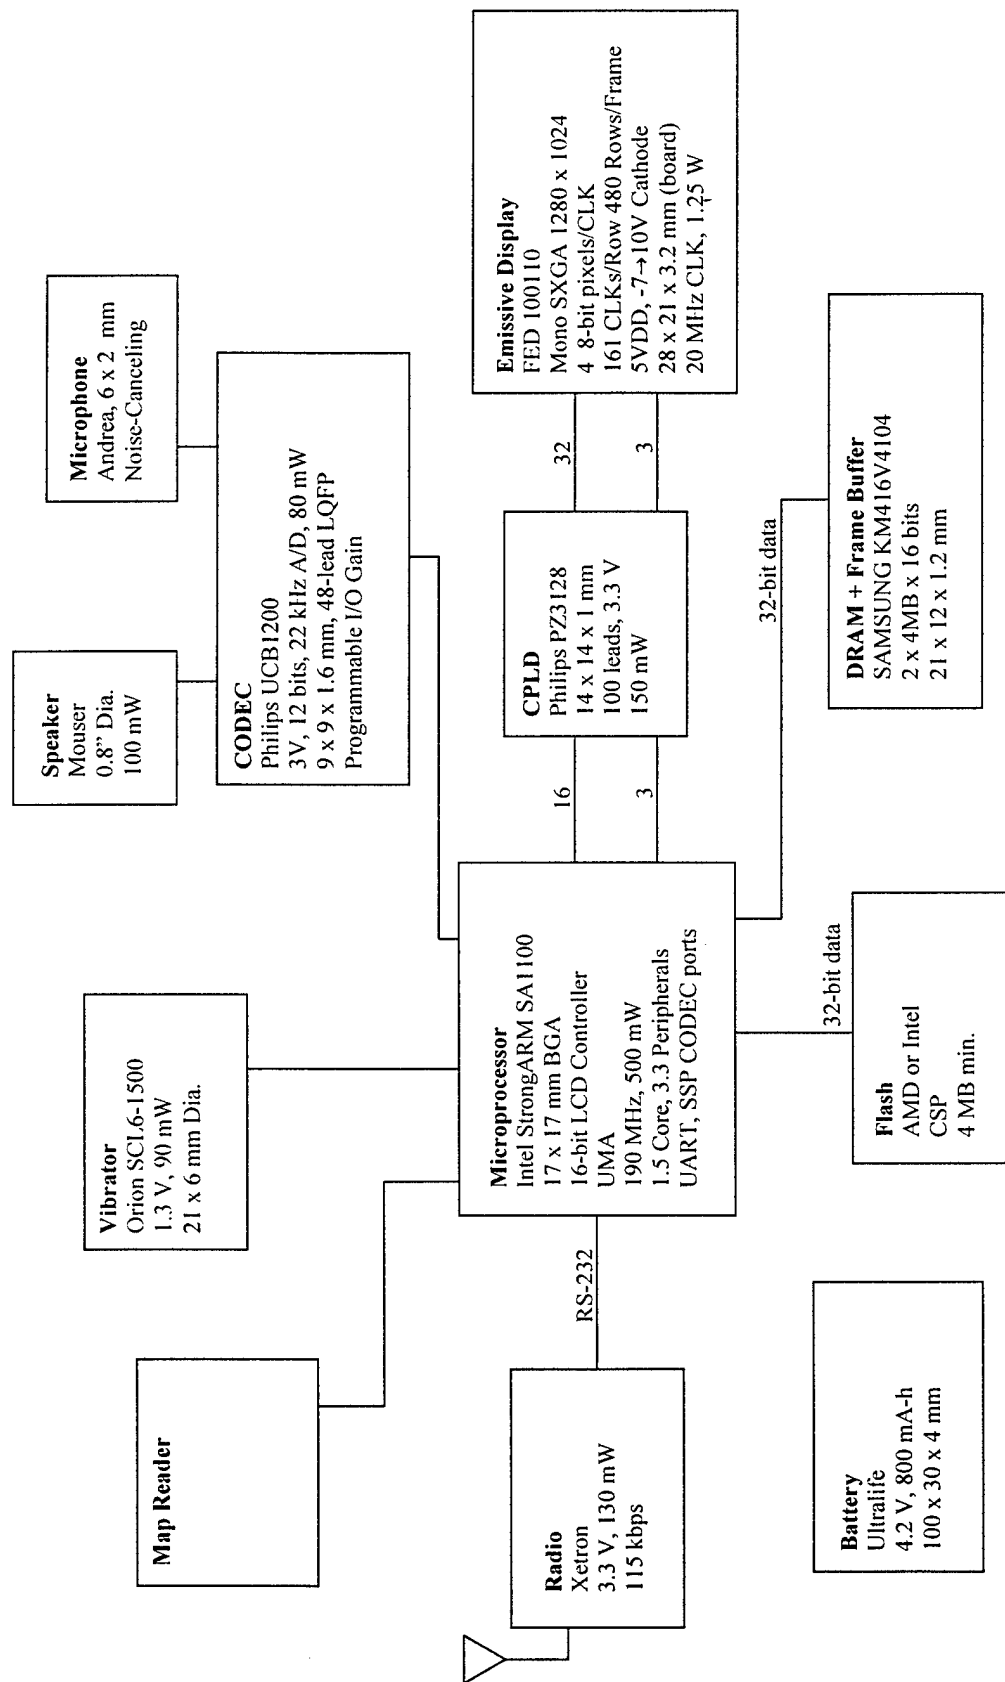


Fig. 2.4 WID Block Diagram With FED Display and Philips CODEC

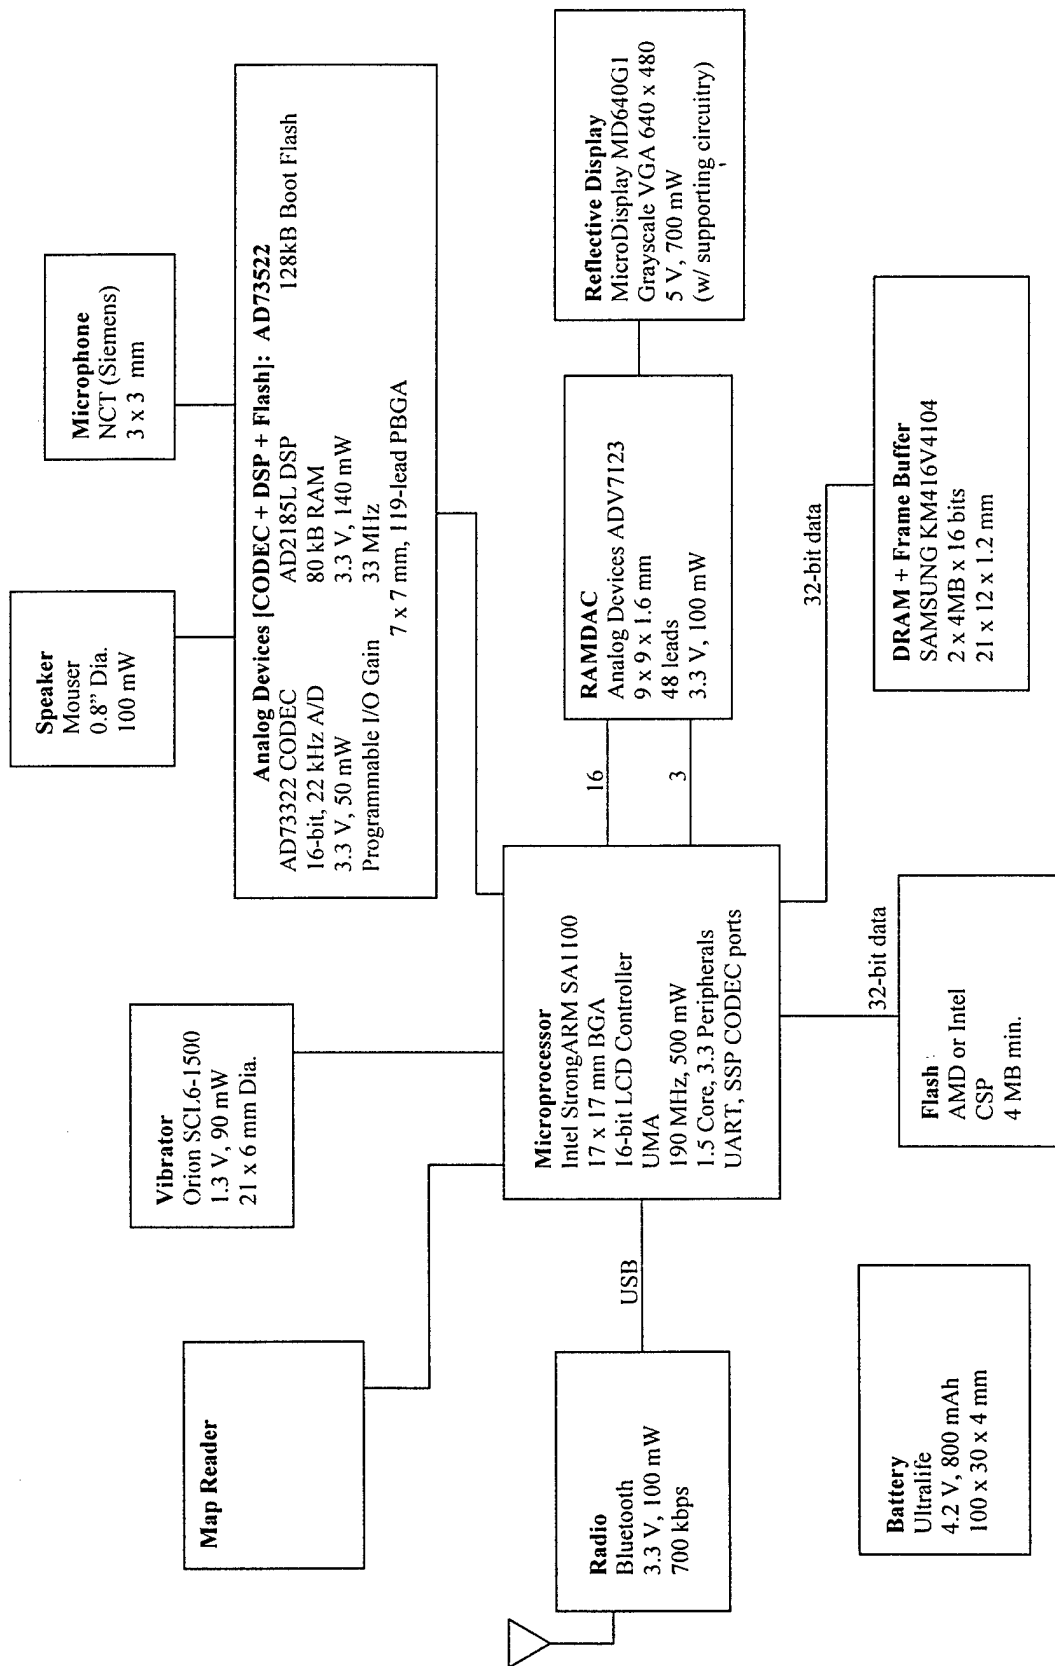


Fig. 2.5 WID Block Diagram with MicroDisplay Display and Analog Devices CODEC/DSP

Chapter III

DEVELOPMENT OF WID EVALUATION BOARD #1

ViA has completed the hardware development of the first evaluation board for the WID. This effort took several months. In this chapter the thought process that led to the configuration of Board 1 is retraced and the engineering tradeoffs made along the way are discussed. The board schematics are then presented. It should be emphasized that the main purpose of this first board is to evaluate different subsystems and components of the WID, not to minimize the form factor of the electronics.

Comparison of the StrongARM SA1100 and NEC Vr4111 Microprocessors

This section addresses the issue of deciding which RISC microprocessor is best suited to the WID platform. Because the specification of the WID platform is still evolving, a broad range of processor features and characteristics necessarily needs to be examined. The WID architecture is impacted principally by the following factors:

- choice of display
- power budget
- space available for supporting chips and circuitry

Different displays have different interfaces to the host processor and may or may not require a frame buffer with associated video controller. The maximum admissible power consumption for the WID is somewhere between 1-2 W. The processor plus all the other components must fit within this limit. Since space inside the WID housing or on the wristband is severely limited, the greater the integration of the electronics the better. Finally, since the bandwidth of the RF interface is limited to a few hundred kbps, the microprocessor needs to be able to interpret fairly high-level Windows commands to drive the display. This means that a fast processor is required.

Among the microprocessors available today or in the near-term, this section examines 32 and 64-bit CPUs made by Intel and NEC, since they utilize the most advanced 0.25μ fabrication process. The processors to be discussed are classified in terms of the following parameters:

- Level of integration:
LCD controller, Audio CODEC, DSP, RAM/ROM, Cache, MAC
- Number and types of interfaces:
Video controller, Audio (A/D specs), Serial (bit rate)
- Efficiency:
Power consumption, Speed, MIPS, MIPS/Watt
- Physical size and packaging

The Motorola DSP56662 processor was dropped from consideration. This decision was based on its lower clock speed and older fabrication process, and by the preference for an integrated LCD controller (SA1100) or CODEC (Vr4111) over the Motorola processor's DSP. Selected parts of what follows were taken directly from the technical reference manuals of the processors they describe.

bus (APB) where all other peripherals are located because it is a higher speed synchronous bus that is able to maintain the data rate required for demanding displays, such as dual-panel color. The LCD's DMA contains two channels that transfer data from external memory to the input FIFO. One channel is used for single-panel displays and two are used for dual-panel displays. The LCD controller outputs to 8 data and 4 control pins. GPIO[9:2] are additionally used for 16-bit/pixel active-mode operation.

Serial Port 4 (MCP/SSP). Serial port 4 contains two separate full-duplex synchronous serial interfaces. The multimedia communications port (MCP) provides an interface to the Philips UCB1200 CODEC. This device has an audio CODEC, a telecom CODEC, a touch-screen interface, four general-purpose analog-to-digital converter inputs, and ten programmable digital I/O lines. The MCP interface is used by the SA1100 both to input and output digital data to and from the CODEC, and to configure and acquire status information from the CODEC's sixteen registers. The synchronous serial port (SSP) is used to interface to a variety of analog-to-digital converters, audio and telecom CODECs, memory chips, and keypad controllers as well as other miscellaneous serial devices. The SSP supports the National Microwire and Texas Instruments synchronous serial protocols as well as a subset of the Motorola serial peripheral interface (SPI) protocol.

Serial Port 3 (UART Port). This port can be used as a debug port.

Serial Port 2 (IrDA Port). This port can be hooked up to the map reader.

Serial Port 1 (AppleTalk). This port can be used for the Xetron radio.

Serial Port 0 (USB Port). This port could not be used in the future to interface to a faster radio chipset because it is a slave.

Efficiency

The SA1100 uses 300 mW at 190 MHz with a core voltage (VDD) of 1.5 V and an external voltage (VDDX) of 3.3 V. It is important to realize that this power figure includes the LCD controller and its accessing of the frame buffer. The power required to maintain the DRAM data is very small compared to the power required of the processor to access the data, and is therefore neglected in the WID power budget estimation.

Packaging

The SA1100 is currently available in a 17 mm x 17 mm BGA package.

NEC Vr4111

NEC's most integrated microprocessor chip is the Vr4181. Unfortunately, this chip targets platforms with lower functionality than the WID: the embedded CODEC does not have sufficiently high resolution and the embedded LCD controller can only drive a 1/4-VGA screen. Thus, the chip set Vr4111 + Vrc4171 becomes the next best alternative. Fig. 3.2 shows the block diagram of the Vr4111 64-bit processor. Although this processor has an embedded CODEC, it does not have sufficiently high resolution and programmable gain control, so for both the

SA1100 and the NEC chip set the same external CODEC (AD73311) would have to be used. A good feature that the Vr4111 has and that the SA1100 does not is a Multiply and Accumulate processor function. The SA1100 achieves the same functionality with more cycles, so this decreases the effective clock speed of the StrongARM relative to the NEC part.

The 64-bit core and the MAC function of the Vr4111 make the relative difference in clock speeds between these two processors smaller than it appears. Even so, it seems doubtful that the NEC part would have an effective speed equal to or faster than the StrongARM's. The combination of the need for two chips and the slower clock frequency make the NEC chip set less desirable. In addition, while the SA1100 uses Unified Memory Architecture, the NEC chips do not. In other words, the fact that the SA1100's LCD controller is on the system bus means that the CPU and the LCD controller can access the same DRAM memory chip. The NEC parts, on the other hand, would require two separate chips, one for the system RAM and another for the frame buffer, which takes up more room. Even though the SA1100 currently supports only EDO RAM, which is an older technology, by next year it will support SDRAM like the Vr4111.

Based on these facts, the SA1100 is considered a more suitable microprocessor with which to start the development of the WID.

Table 3.1 Comparison of SA1100 and Vr4111 Microprocessors			
	SA1100	NEC Vr4111	NEC Vrc4171
	Microprocessor	Microprocessor	Video Controller
Level of Integration			
LCD Controller	Yes	No	Yes
Audio CODEC	No	Yes	No
DSP	No	No	No
RAM/ROM	No	No	No
Cache	16k, 8D	16k, 8D	No
DMA controller	Yes	Yes	No
MAC (multiply and accumulate)	No	Yes	No
Number and Types of Interfaces			
Video Controller	Digital (16)	-	Digital (16)
Audio	No	10-bit A/D/A	No
Serial UART	16550 250 kbps	Yes (RS-232)	No
SDLC	230 kbps	No	No
USB	Yes (slave)	No	No
Efficiency			
Power Consumption (mW)	300	250	225
Speed (MHz)	190	100	-
MIPS	230	130	-
MIPS/Watt	460	274	
Physical			
Core Voltage	1.5V	2.5V	3.3 V
Process	.35u	.25u	.25u
Size	17mm x 17mm		28 mm x 28 mm
Packaging	MBGA	224 FPBGA	208 LQFP




Fig. 3.2 Block Diagram of the Vr4111 Processor

Data Circuit

The FED display presents an unusual interface to the host processor. In order to refresh its 1280 x 1024 pixels at a sufficiently fast rate without requiring an excessively fast system (pixel) clock, its IC logic is configured to paint 4 pixels per pixel clock, in a row. When set in VGA mode, each of the four VGA pixels provided on the input pins is automatically duplicated four times in a square array, thereby generating a 2 x 8 array of SXVGA pixels. In this mode, therefore, only 160 pixel clocks per line and 480 lines per frame are needed.

Although the SA1100 LCD controller reads 32 bits at a time from the frame buffer, the data are handled differently depending on the choice of bits/pixel programmed with the parameter PBS (Pixel Bit Size). For 4 or 8 bits/pixel, the LCD controller assumes them to be an index into a palette stored at the top of the frame buffer, in a separate area of the system's main memory. It also assumes a passive-mode screen and feeds the palette data to the dither engine (internal to the LCD controller). 12 bits/pixel are assumed to imply color: in passive mode the data bypass the palette but are still sent to the dither engine, while in active mode they index the palette and bypass the dither engine. Finally, 16-bits/pixel mode bypasses both the palette and the dither engine. Fig. 3.4 summarizes these options.

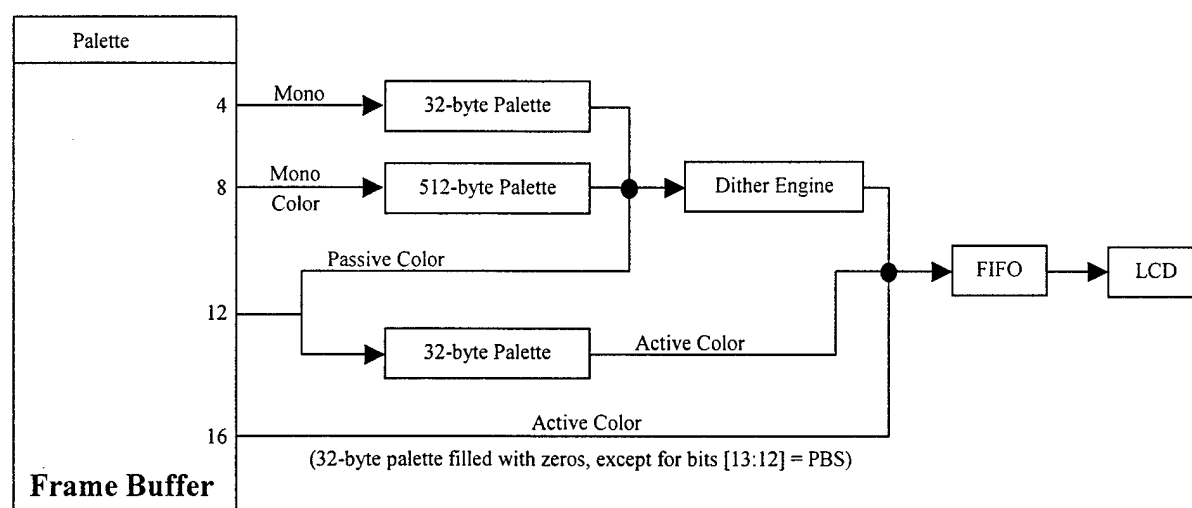


Fig. 3.4 Data paths of SA1100 LCD controller for different pixel bit sizes (PBS)

Fig. 3.5 shows a schematic of the data circuit and the 32-bit interface of the FED display. The mode that can be made compatible with the FED display most easily is active-color, 16 bits/pixel. Since 16 bits/pixel is the largest word that the LCD controller can provide on its output pins, some glue logic is needed between the SA1100 and the FED display. Each LCD controller clock cycle (L_PCLK), a PAL (Programmable Array Logic) chip containing a 32-bit buffer can be used to hold the current 16 bits, which are written high or low every other cycle. Every two L_PCLK cycles (= 1 FED CLK cycle) the 32 bits are output by the PAL onto the FED input pins. Here a DAC converts the bit information into analog voltages that charge the pixel capacitors. The overall frame refresh rate will depend on how long the pixel capacitors can hold the charge. As of this writing, FED is still measuring this quantity.

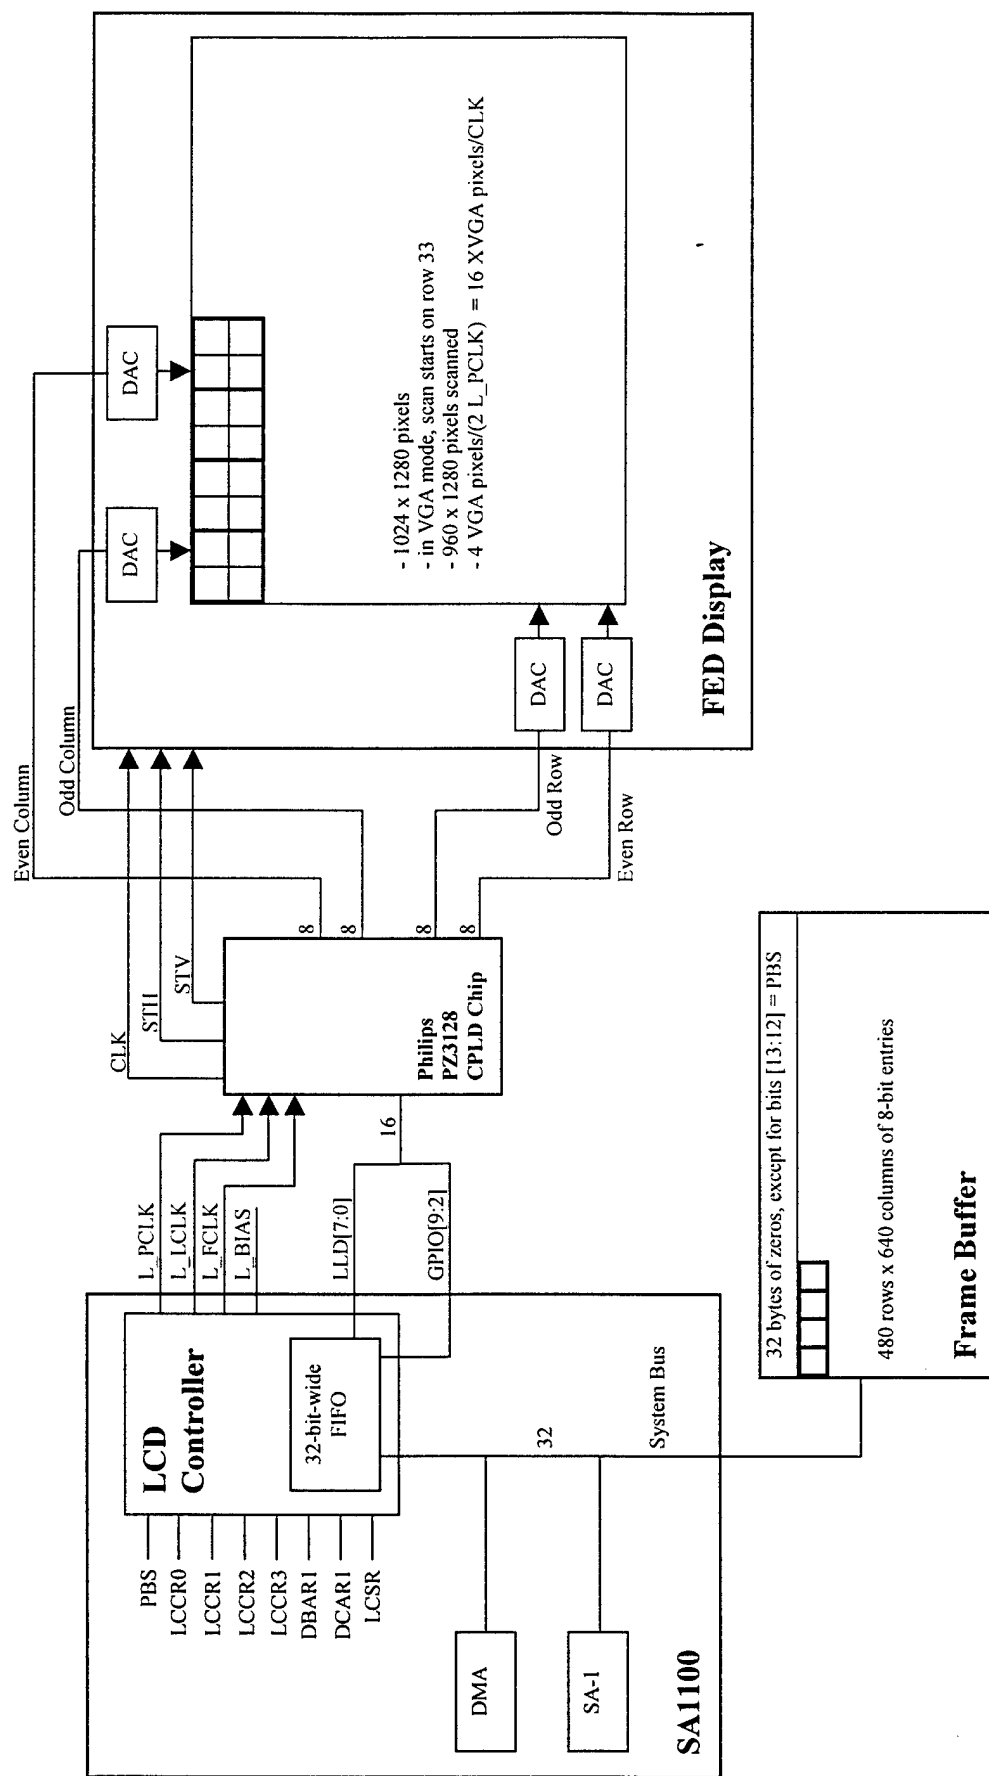


Fig. 3.5 Block diagram of data flow from the SA1100 to the FED display.

LCD Control Registers

Before discussing the switching circuit, it is helpful to clarify what inputs are needed to program the LCD controller to conform to different displays. These inputs are in the form of bit fields in control registers and memory locations, shown in Fig. 3.5. For each register, the first table, below, lists the names and a brief description of the bit fields, and gives their possible values. At the bottom of each table the location and values of the bit fields in the registers are shown. The fields left blank have more than one setting compatible with the FED display, and their values will be determined at a later date. The values given represent settings that are compatible with the FED display. While the data stream can be made entirely compatible, the timing signals are all compatible except for one, to be discussed more fully below.

Table 3.2 First 16 Bits of Palette in Frame Buffer

Table 3.2 First 16 Bits of Palette in Frame Buffer															
PBS		Pixel Bit Size 00 – 4 bits/pixel 01 – 8 bits/pixel 10 – 12 or 16 bits/pixel 11 – reserved												2 bits	
15	14	PBS		11	10	9	8	7	6	5	4	3	2	1	0
0	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0

Table 3.3 Control Register 1 – LCCR0

Table 3.3 Control Register 1 – LCCR0															
LEN		LCD Enable 0 – disabled 1 – enabled											1 bit		
CMS		Color/Mono 0 – color 1 – monochrome											1 bit		
SDS		Single/Dual Panel 0 – single 1 – dual											1 bit		
PAS		Passive/Active 0 – passive 1 – active											1 bit		
BLE		Big/Little Endian 0 – little 1 – big											1 bit		
DPD		Double-Pixel Data (not relevant to active mode) 0 – LDD[3:0] 1 – LDD[7:0]											1 bit		
FDD		FIFO DMA Request Delay 0 → 255 memory control clocks to wait (1 mem clock = 2 pixel clocks)											8 bits		
FDD				11	10	DPD	BLE	PAS	6	5	4	3	SDS	CMS	LEN
				-	-	0	0	1	-	-	-	-	0	0	1
31	30	29	28	27	26	25	24	23	22	21	20	FDD			
-	-	-	-	-	-	-	-	-	-	-	-				

Table 3.4 Control Register 2 – LCCR1

Table 3.4 Control Register 2 – LCCR1															
PPL	Pixels Per Line 640														10 bits
HSW	Horizontal Sync Pulse Width 0 → 63 <i>pixel</i> clock periods														6 bits
ELW	End-of-Line Pixel Clock Wait Count 0 → 255 <i>pixel</i> clock periods														8 bits
BLW	Beginning-of-Line Pixel Clock Wait Count 0 → 255 <i>pixel</i> clock periods														8 bits
HSW						PPL									
						1	0	1	0	0	0	0	0	0	0
BLW								ELW							

Table 3.5 Control Register 3 – LCCR2

Table 3.5 Control Register 3 – LCCR2															
LPS	Lines Per Screen 480														10 bits
VSW	Vertical Sync Pulse Width 0 → 63 <i>line</i> clock periods														6 bits
EFW	End-of-Frame Line Clock Wait Count 0 → 255 <i>line</i> clock periods														8 bits
BFW	Beginning-of-Frame Line Clock Wait Count 0 → 255 <i>line</i> clock periods														8 bits
VSW						LPS									
						0	1	1	1	1	0	0	0	0	0
BFW								EFW							

Table 3.6 Control Register 4 – LCCR3

PCD	Pixel Clock Divider 0 → 255 $\text{Pixel Frequency} = \frac{\text{CPU Frequency}}{2(\text{PCD} + 2)}$	8 bits
ACB	AC Bias Pin Frequency 0 → 255 no effect in active mode	8 bits
API	AC Bias Pin Transitions Per Interrupt 0 → 15 see reference manual	4 bits
VSP	Vertical Sync Polarity 0 – L_FCLK active high 1 – L_FCLK active low	1 bit
HSP	Horizontal Sync Polarity 0 – L_LCLK active high 1 – L_LCLK active low	1 bit

Switching Circuit

The switching circuit of the FED display requires only three external timing signals: pixel clock (CLK), horizontal (line) clock (STH), and vertical (frame) clock (STV). The signals themselves are pulses of a specified width. The FED display requires the approximate values shown in Table 3.10. Note that FED has not released their final spec sheet, so the values discussed may change.

Fig. 3.6 shows how the timing signals from the SA1100 are interfaced to the switching circuit of the FED display. Table 3.10 shows values that can be chosen for the SA1100 pixel clock (through the SA1100 parameter PCD) and for other SA1100 parameters that make all but one interface value compatible with the FED display. The only signal that remains incompatible with the FED display is the pulse width of the vertical signal clock. While the SA1100 follows the general norm of making this pulse width equal to one or more line periods, the FED requires this pulse width to be exactly one system clock period ($1/\text{CLK}$). The reason is that the internal IC of the FED display uses the *level* (rather than the left edge) of the vertical synchronization signal to reset itself for a new frame. If this signal stays high for more than one system clock, the FED display will keep resetting itself. Thus, some glue logic will be needed to ensure that the high level of the L_FCLK pulse provided by the SA1100 at the beginning of each frame is brought to low after one system clock period. This check will be programmed into the PAL chip shown in Fig. 3.6. At the present time, for this chip ViA is planning to use the Philips PZ3128 CPLD chip.

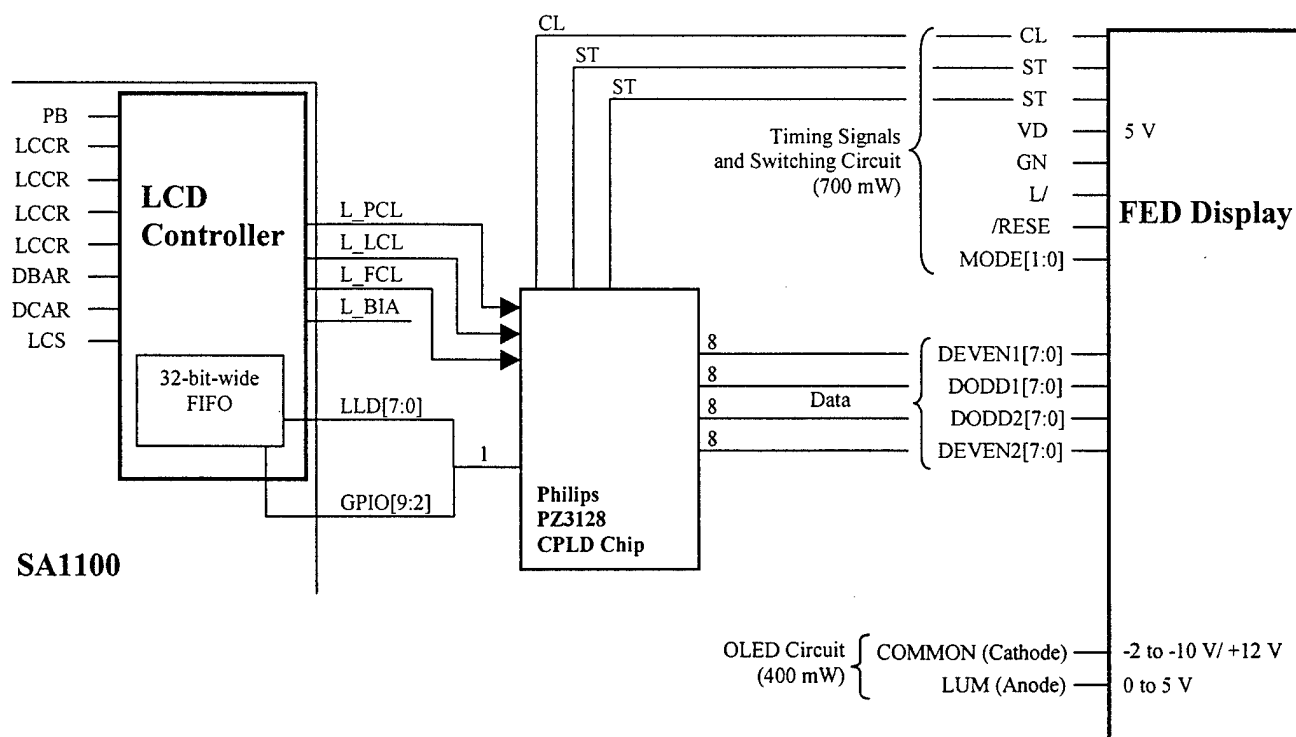


Fig. 3.6 Block Diagram of the FED display interface

Table 3.10 Worksheet for calculating the timing parameters required to interface the SA1100 to the FED display

SA1100 Symbol	FED Symbol	Description	SA1100 Range		FED Range		SA1100 Parameter		FED Value Chosen	Unit
			Min	Max	Min	Max	CPU			
L_PCLK	Fclk	System (Pixel) Clock Frequency	0.37	31.67	10	40	PCD	2	23.75	MHz
	Tclk	Pixel Clock Period	32	2705	25	100			42.11	ns
L_LCLK	Tsth	Horizontal Sync Clock Period		(any)	8	32	PPL	Tclk * 640	26.95	µs
	Twsth	Horizontal Sync Pulse Width	42.11	2652.63	42.11	42.11	HSW	1	42.11	ns
L_FCLK	Tstv	Vertical Sync Clock Period		(any)	8.2	16.6	LPS	Tsth * 480	12.93	ms
	Twstv	Vertical Sync Pulse Width	26947	1697684	42.11	42.11	VSW	1	42.11	ns
	Tf	Frame Period			8	16.66				ms
	Trst	Reset Pulse Width			100	...			,	µs
	Ta	Ambient Temperature			-30	70				deg. C

Interfacing the SA1100 to the MicroDisplay Display

Interfacing the SA1100 to the Microdisplay display is not a straightforward task, mostly because the display is under constant evolution. MicroDisplay has the capability to integrate all the interfacing circuitry onto the display chip, but will not necessarily have this task completed in a useful timeframe for the WID development. Based on the information released to ViA, MicroDisplay has the following versions of their display completed or planned:

Table 3.11 Summary of MicroDisplay's Display Development			
Product/Display Type	Platform	Power (mW)	Total Power (mW)
1 - B&W, Board-level VGA interface/decoder	2/3 credit-card size board	1500	1500
2 - B&W, Chip-level VGA interface/decoder	Display chip & decoder	60	280
• samples available	LED Illuminator	20	
• volume production pending Generation II design	PLL	100	
	RAMDAC	100	
3 - Field-sequential color display (board-level)	System Development Board	4200	4200
4 - Field-sequential color display (chip-level)	Display chip	30	1100
• has memory-mapped and VGA interface	RGB LED Illuminator	80	
	ASIC + SRAM FB	1000	
5 - Diffraction-gratings color display	Not ready for production	TBD	TBD
6 - SVGA-resolution field-sequential color		TBD	TBD
• samples available			
• volume production available 4/23/99			
7 - XGA-resolution field-sequential color		TBD	TBD
• samples available			

Out of these possibilities, the low power requirements and high integration of Display 2 are very attractive for the WID. Unfortunately, the fact that it is monochrome is a drawback. In addition, this display currently it is not in production due to errors in the first design. A second design and manufacturing run could be arranged by MicroDisplay for a relatively moderate fee.

The 80 mW quoted for the RGB LED illuminators is an estimate made by ViA based on available LED and optics technology and on brightness requirements for outdoor visibility of reflective displays.

If Display 4 is used with the memory-mapped interface, it would save some of the power of the SA1100 since the display data would be written directly to the SRAM frame buffer by the processor, without relying on the integrated LCD controller (TFT interface). Although ViA has not yet conducted any tests of the SA1100 processor to measure its power consumption, the system power savings brought by the use of a memory-mapped interface are estimated to be on the order of 100 mW.

If a display with a VGA interface is used, ViA will utilize the Analog Devices AD7123 RAMDAC since it is a relatively low-power, 3-V part.

RF Interface

ViA is relying on a radio being developed by Xetron as the initial wireless interface to the host computer, the ViA II. The first hardware implementation of the WID evaluation platform, to be discussed more fully below, will utilize the Xetron development board. ViA, however, has joined the Bluetooth consortium and is planning to rely on a Bluetooth radio as the final solution for the WID. This makes sense both from a technical point of view as well as from a marketing and compatibility point of view. ViA is slated to receive early Bluetooth evaluation boards in Q2 99.

A preliminary block diagram of the Bluetooth system is shown in Fig. 3.7. One of the greatest advantages to relying on this radio is that it comes with a well-defined hardware and software interface that will make it significantly easier to implement. In addition, porting the WID to other host computers that support the Bluetooth specification will not require any additional work.

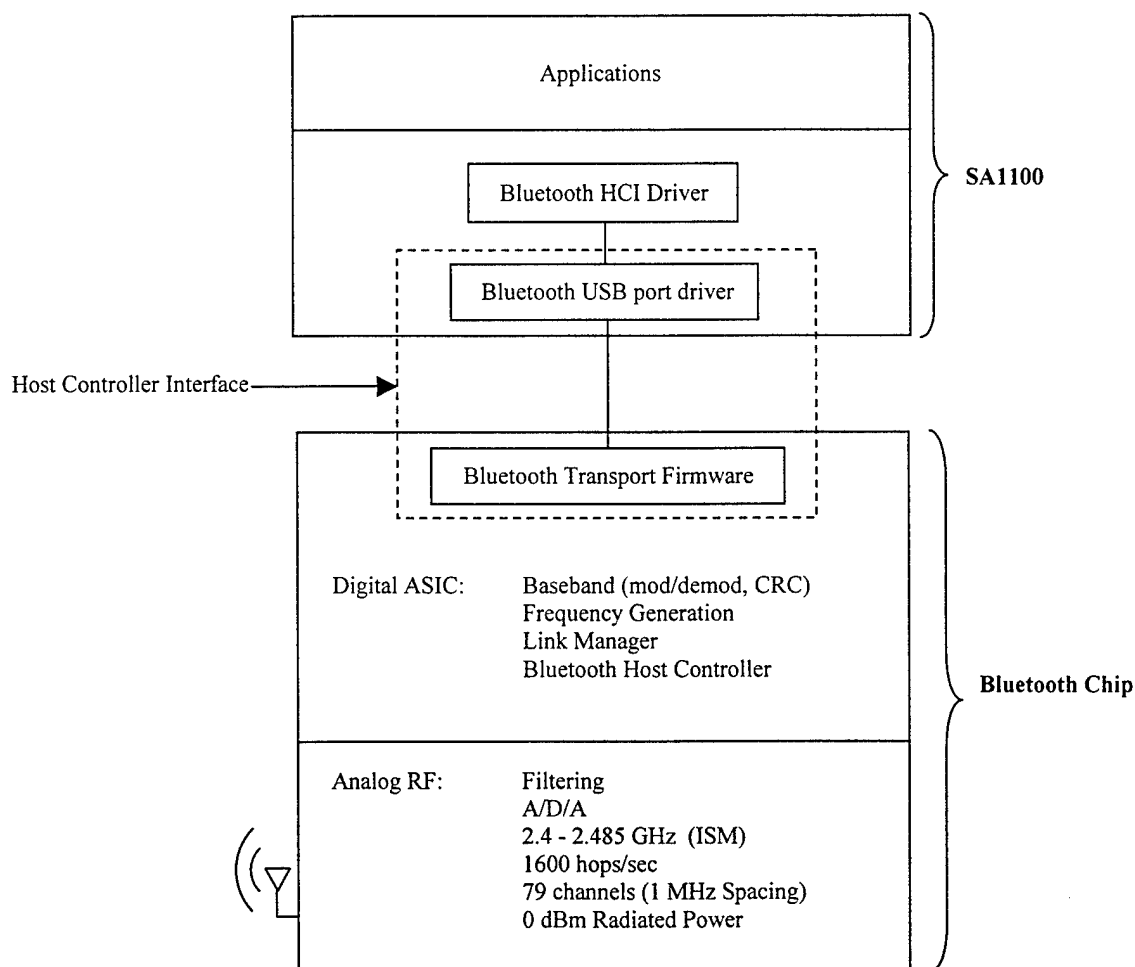



Fig. 3.7 Preliminary Block Diagram of Bluetooth Radio and Drivers

At the time of this writing Xetron evaluation boards have arrived and have been tested. The testing has involved connecting each board to a PC with an RS-232 serial cable and verifying that the two PCs exchange files and data over this link. The transceiver boards can be communicate either over the air or through a serial cable. The RF data rate with the present boards is 19.2 kbps, while with a cable it is 115 kbps. Xetron is planning to ship improved boards to double the RF data rate by the end of March. By this time the first evaluation units of the Bluetooth radio, as implemented by National Semiconductor, should become available. Fig. 3.8 summarizes the functions of the present Xetron evaluation board-level system.



- Board A:** Maxim Semiconductor, 1.2 V - 3.3 V input, DC-DC boost converter.
- Powers all the board from 2 D-Size batteries
- Board B:** Maxim Semiconductor, CMOS RS-232 interface circuit
- Translates TTL-level input signals from the PC to proper RS-232 signal levels and vice versa
- Board C:** Xetron ASIC interface board
- Reads data from RS-232 port
 - Modulates transmitted data
 - Demodulates received data
 - Selects radio channels
 - Controls radiated output power of radio
- Board D:** Xetron Hornet 2.4 GHz transceiver radio.
- Transmits and receives data from ASIC board over the air

Fig. 3.8 Functional Description of Xetron's Hornet Radio Evaluation Board

RAM and Flash

The Intel SideARM evaluation board comes with two Intel 28F016SV-070 (1Mx16) RAM chips. These parts are being phased out and a new chip will have to be used. By Q3 99 Intel will have added support for SDRAM to the SA1100 processor and RAM chips in BGA packaging will be utilized on the WID board being developed at that time.

WID Board #1 Schematics

In the following pages the schematics of the first development board for the WID are presented. As a starting point in the design of the WID hardware, these schematics follow the layout of the Intel SA1100 SideARM development board. To service the functions not supported by the SA1100, Intel has developed a companion chip, the SA1101. Accordingly, Intel has developed the SideKick, a companion development board for the SideARM, to house the SA1101 and related ports. Table 3.12 summarizes the functions performed by this chipset.

Table 3.12 Peripheral Support of StrongARM Chipset	
SA1100	SA1101
LCD Controller	VGA Controller
Serial Port 0: USB (slave)	USB (master)
Serial Port 1: Appletalk (RF)	Parallel Port
Serial Port 2: IrDA (Map Reader)	PCMCIA
Serial Port 3: UART (Debug)	GPIO
Serial Port 4: MCP/SSP (Audio)	Keyboard
	Mouse

The WID board follows a very similar architecture, with the Main Board hosting the SA1100 and the Mezzanine Board hosting the SA1101. Table 3.13 summarizes the functions performed by WID Board #1, and Fig. 3.9 shows the same information in block diagram form.

Table 3.13 Functional Summary of WID Board #1	
Main Board	Mezzanine Board
SA1100 CPU + LCD & DMA Controllers	VGA Controller to CRT Display
RAMDAC to CRT Display	USB (master)
CPLD to Parallel Port	Parallel Port
Serial Port 1: Xetron Radio	Keyboard
Serial Port 2: Map Reader	Mouse
Serial Port 3: Debug	
Serial Port 4: CODEC	

The following figures contain the schematics of the Main and Mezzanine Boards.






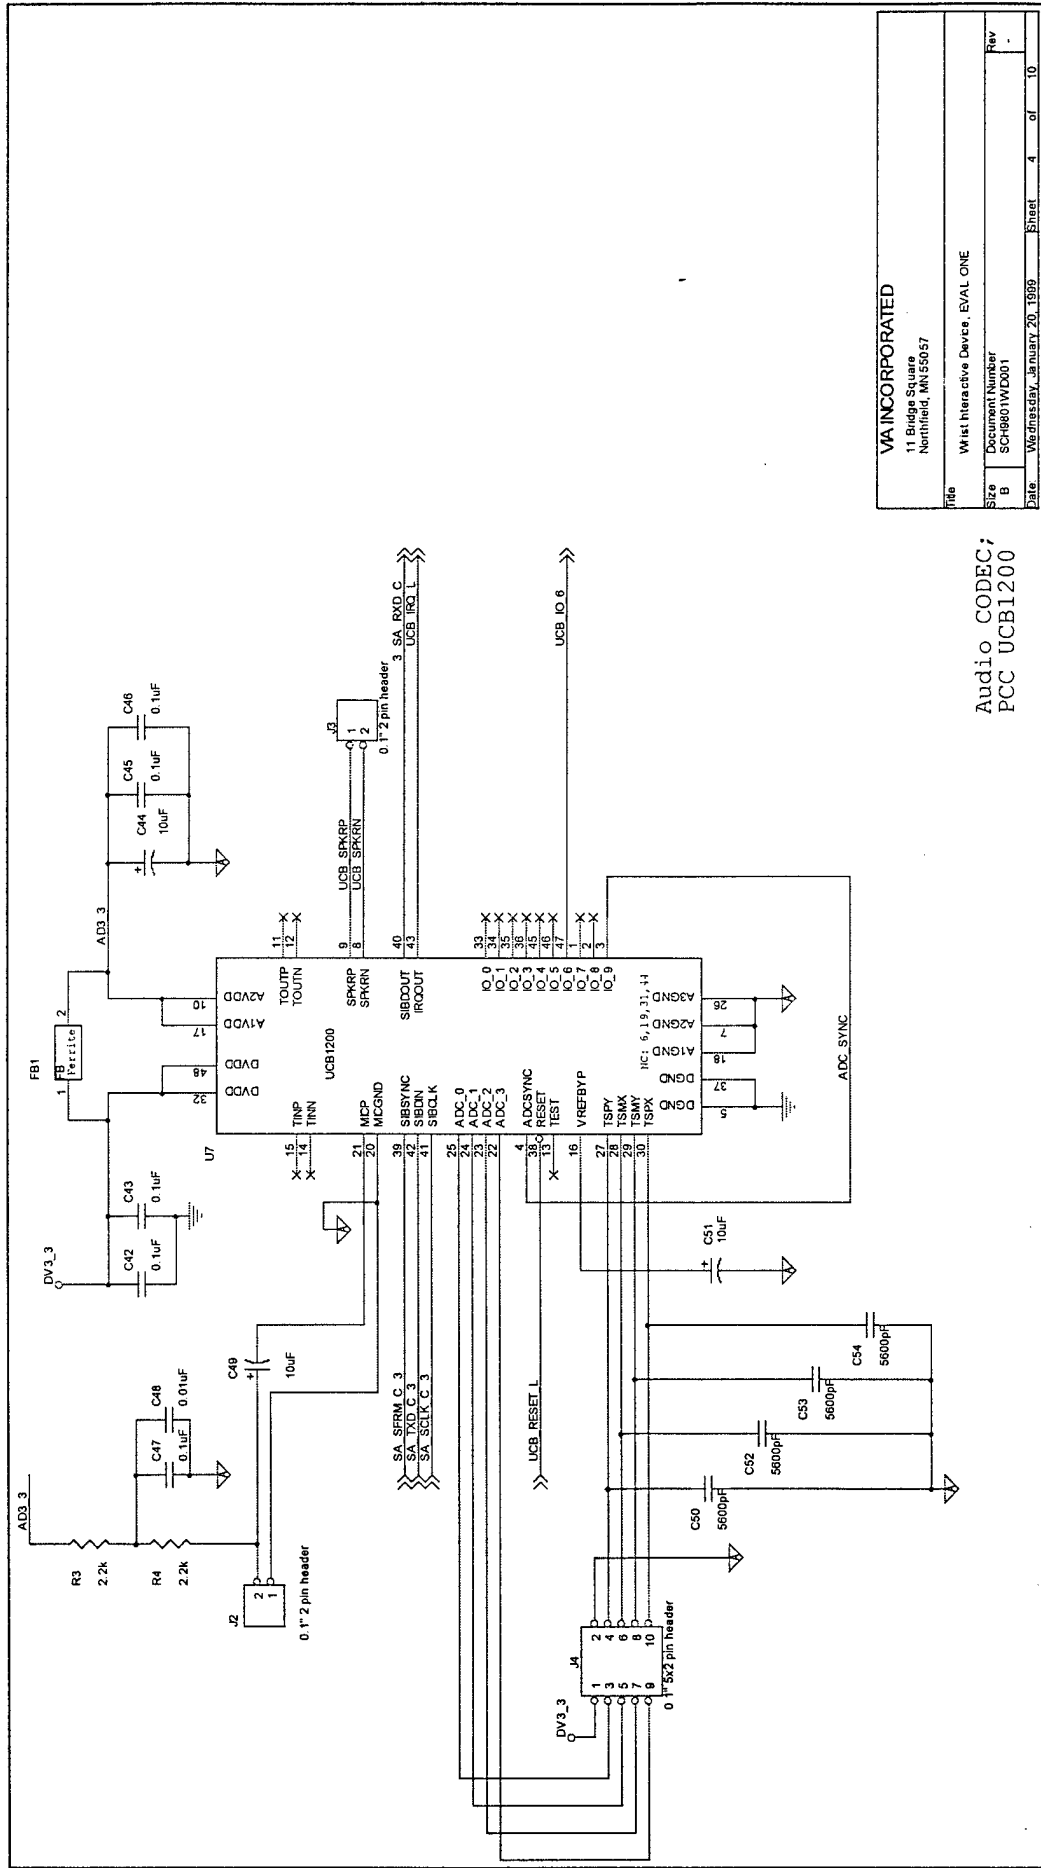
Fig. 3.9 Block Diagram of WID Board #1



Block Diagram

VIA INCORPORATED 11 Bridge Square Northfield, MN 55057		
Title: Wrist Interactive Device, EVAL ONE		
Size: B	Document Number: SCH9801WD001	Rev: -
Date: Wednesday, January 20, 1999	Sheet: 2	of 10





VA INCORPORATED
11 Bridge Square
Northfield, MN 55057

File: Wist Interactive Device, EVAL ONE

Size: Document Number

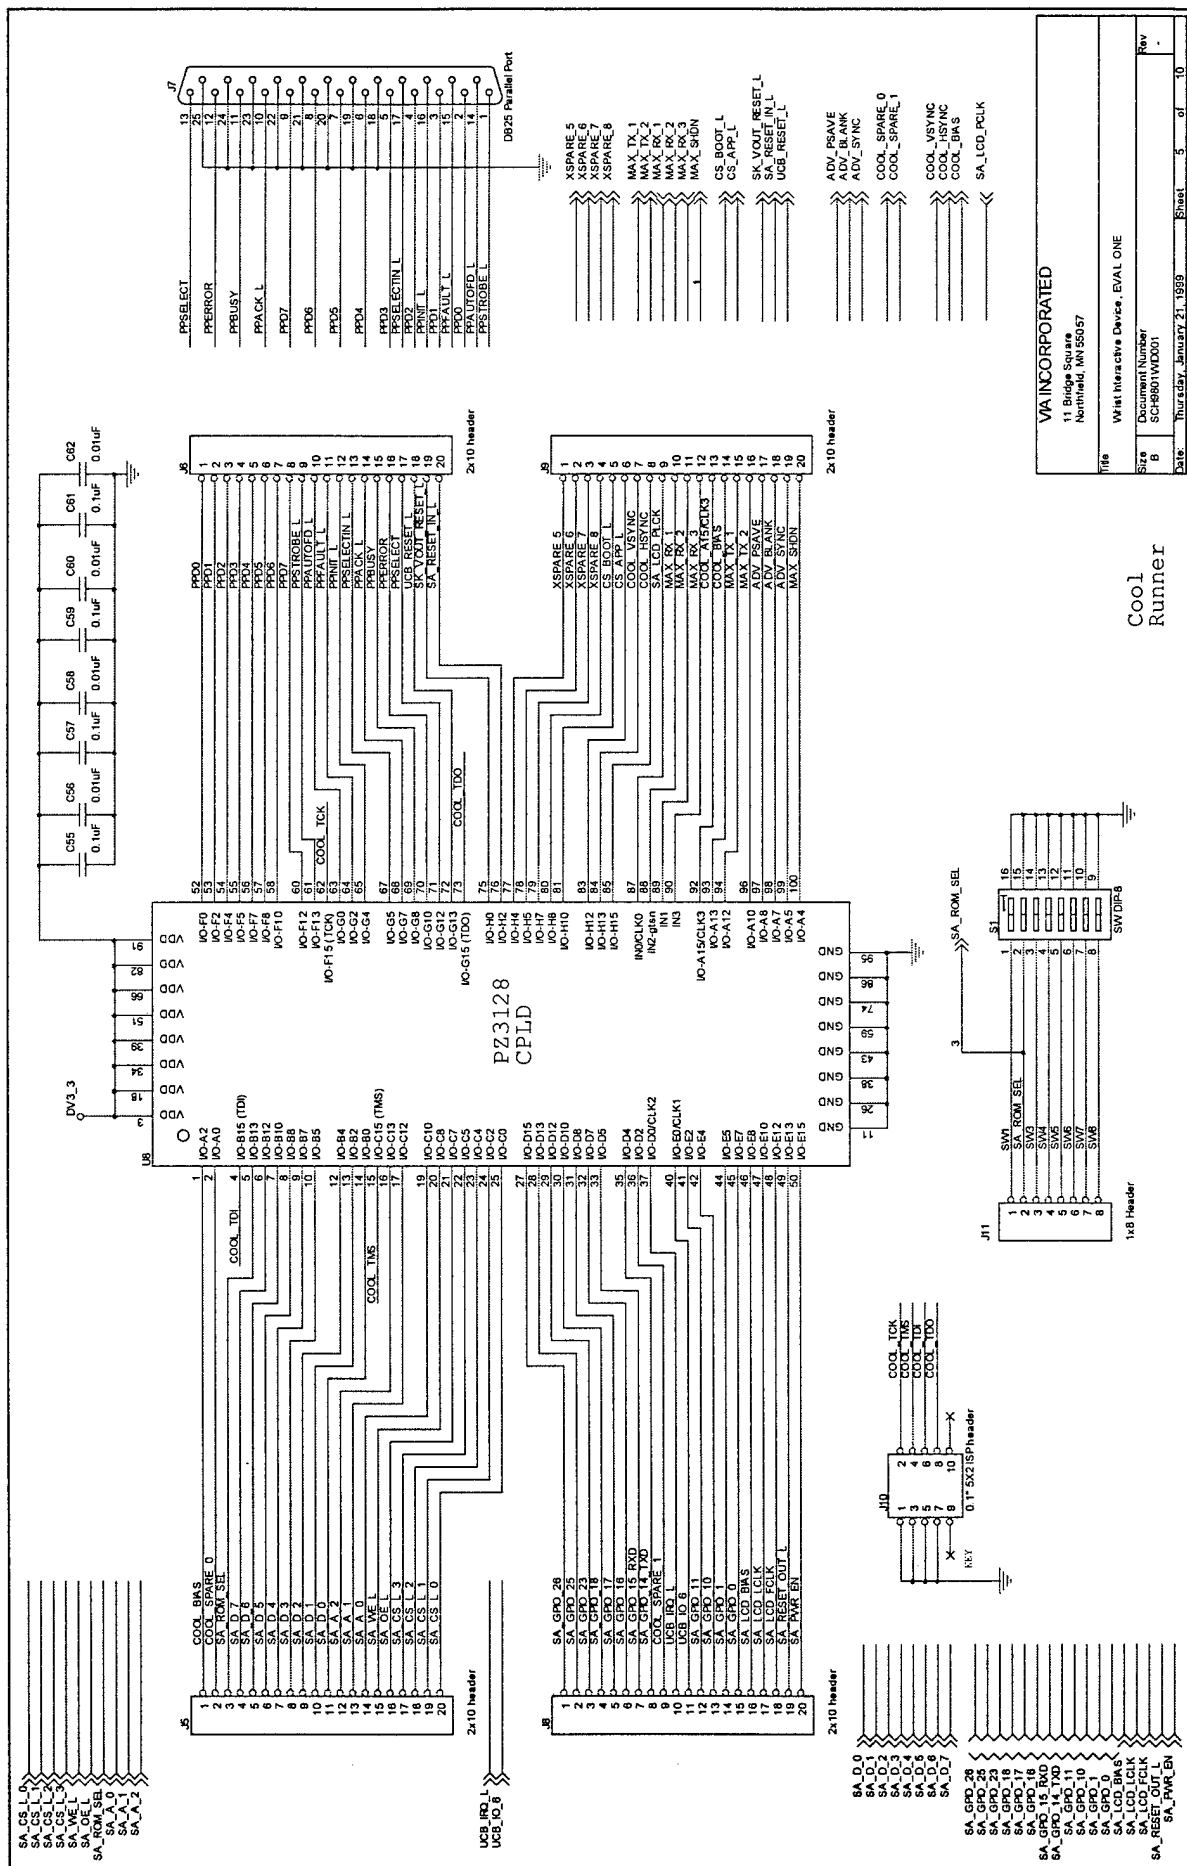
B: SCH9801WD001

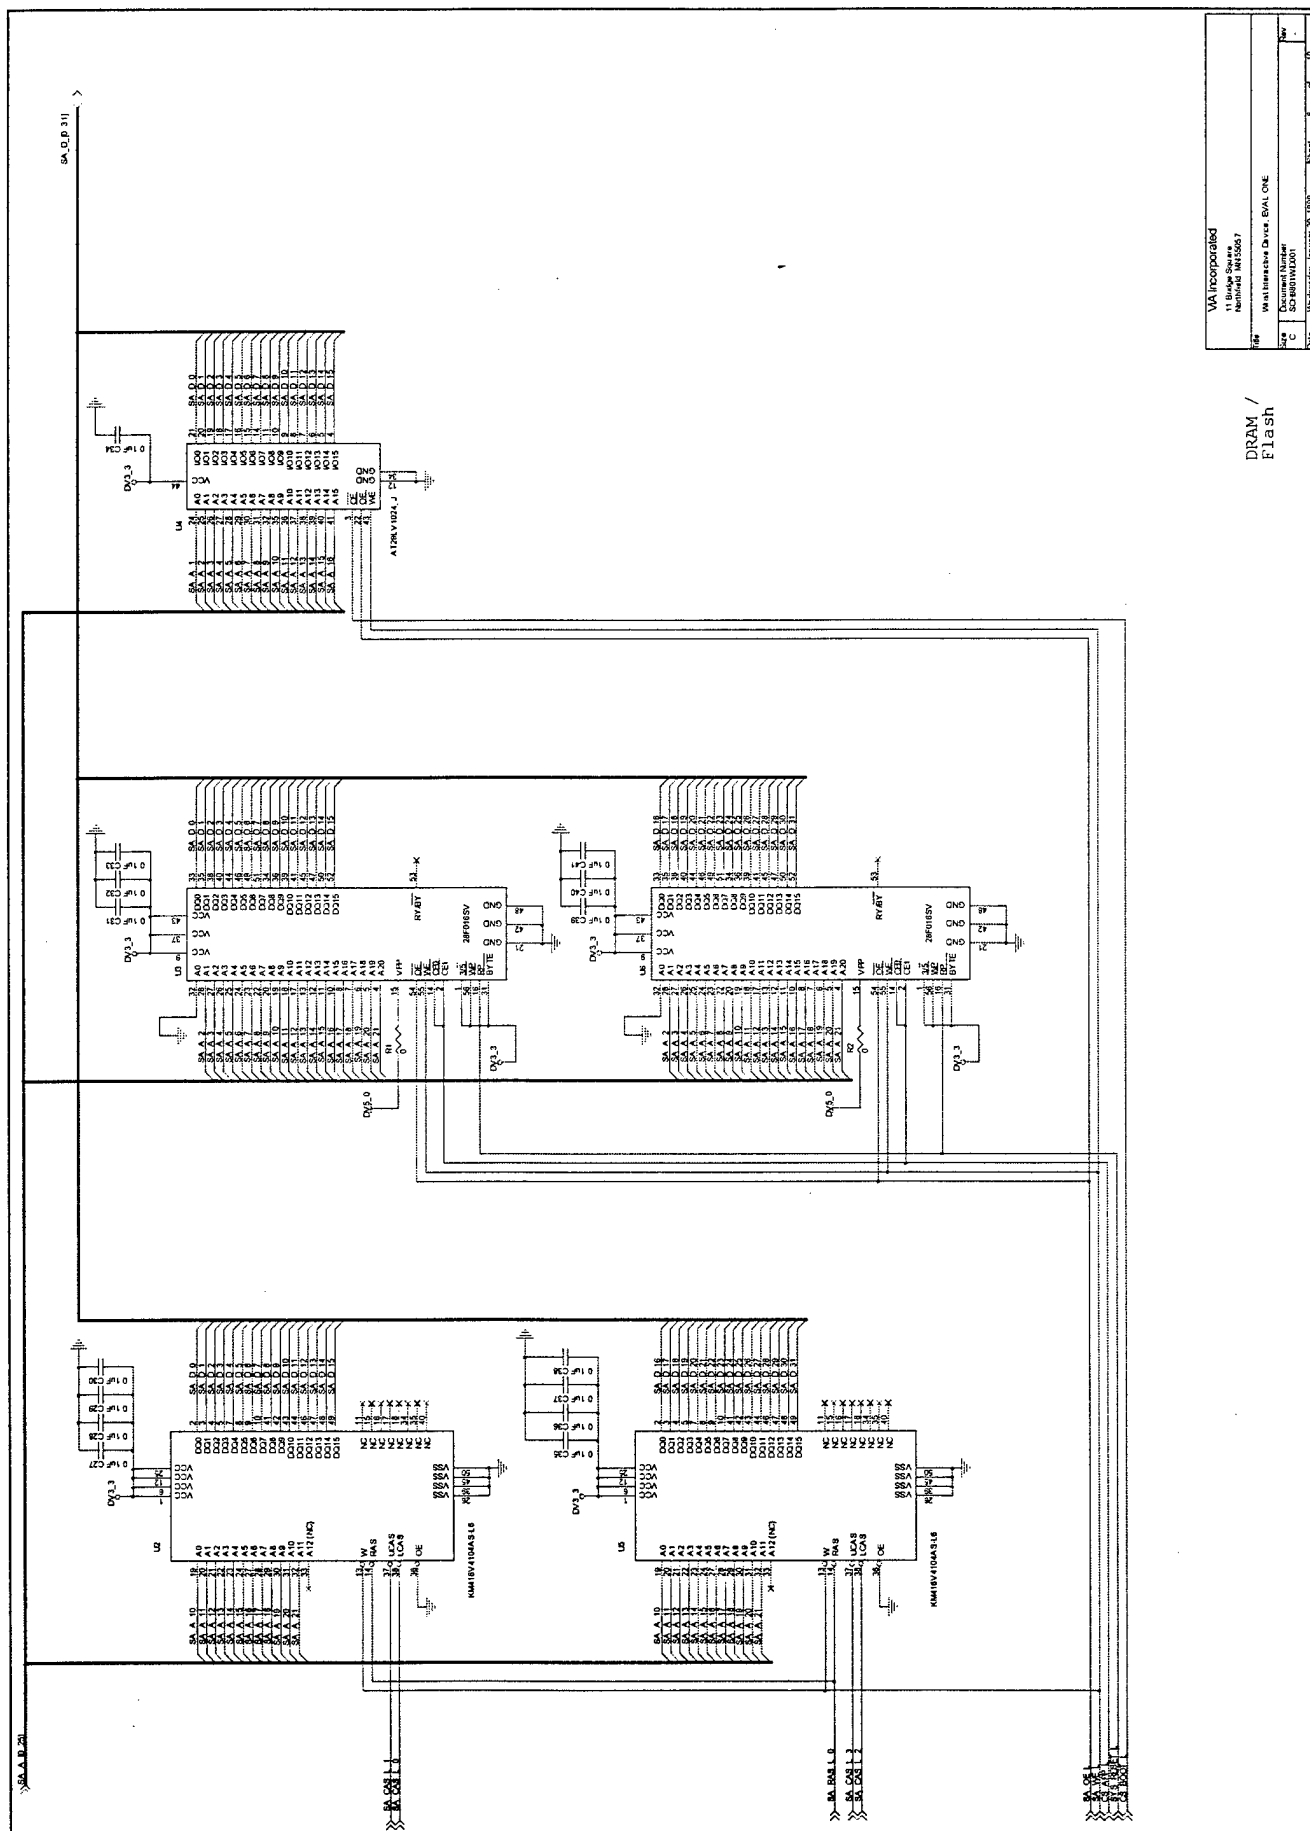
Rev: -

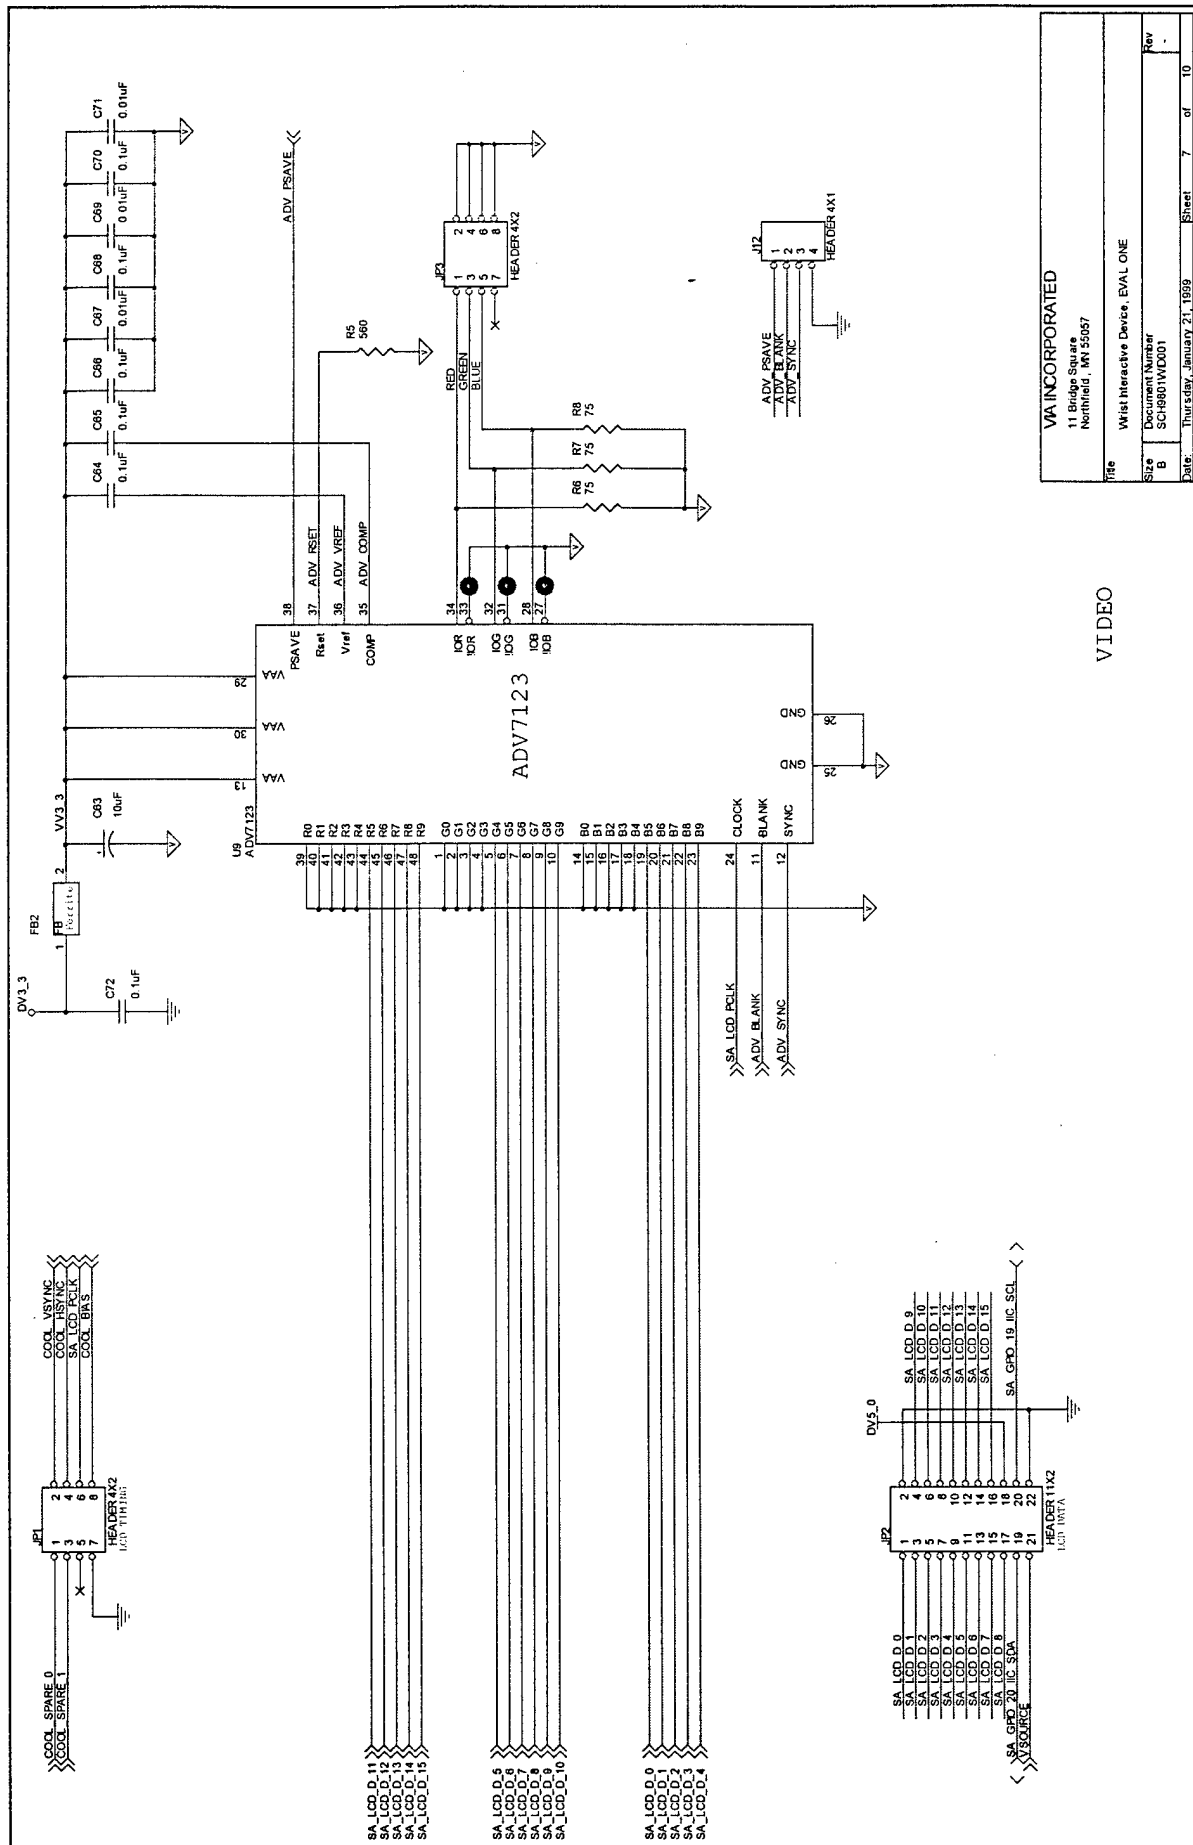
Date: Wednesday, January 20, 1999

Sheet: 4 of 10

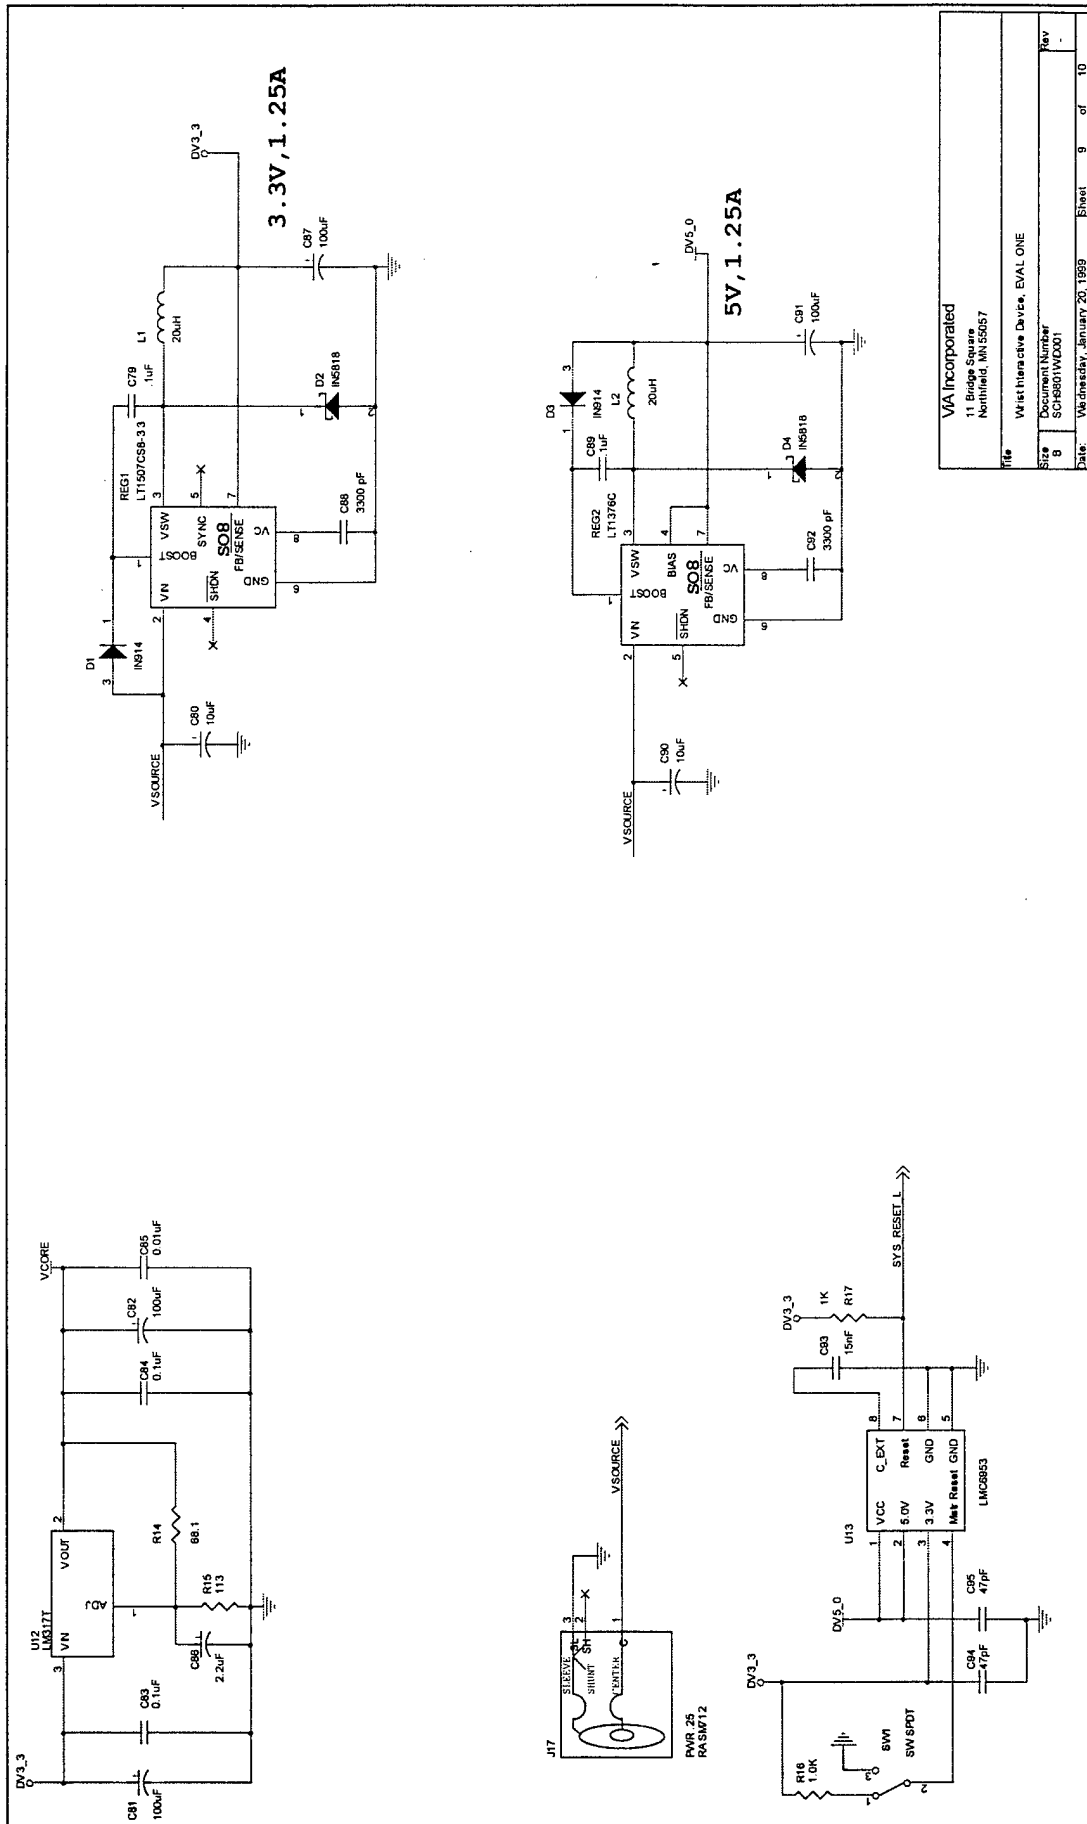
Audio CODEC;
PCC UCB1200



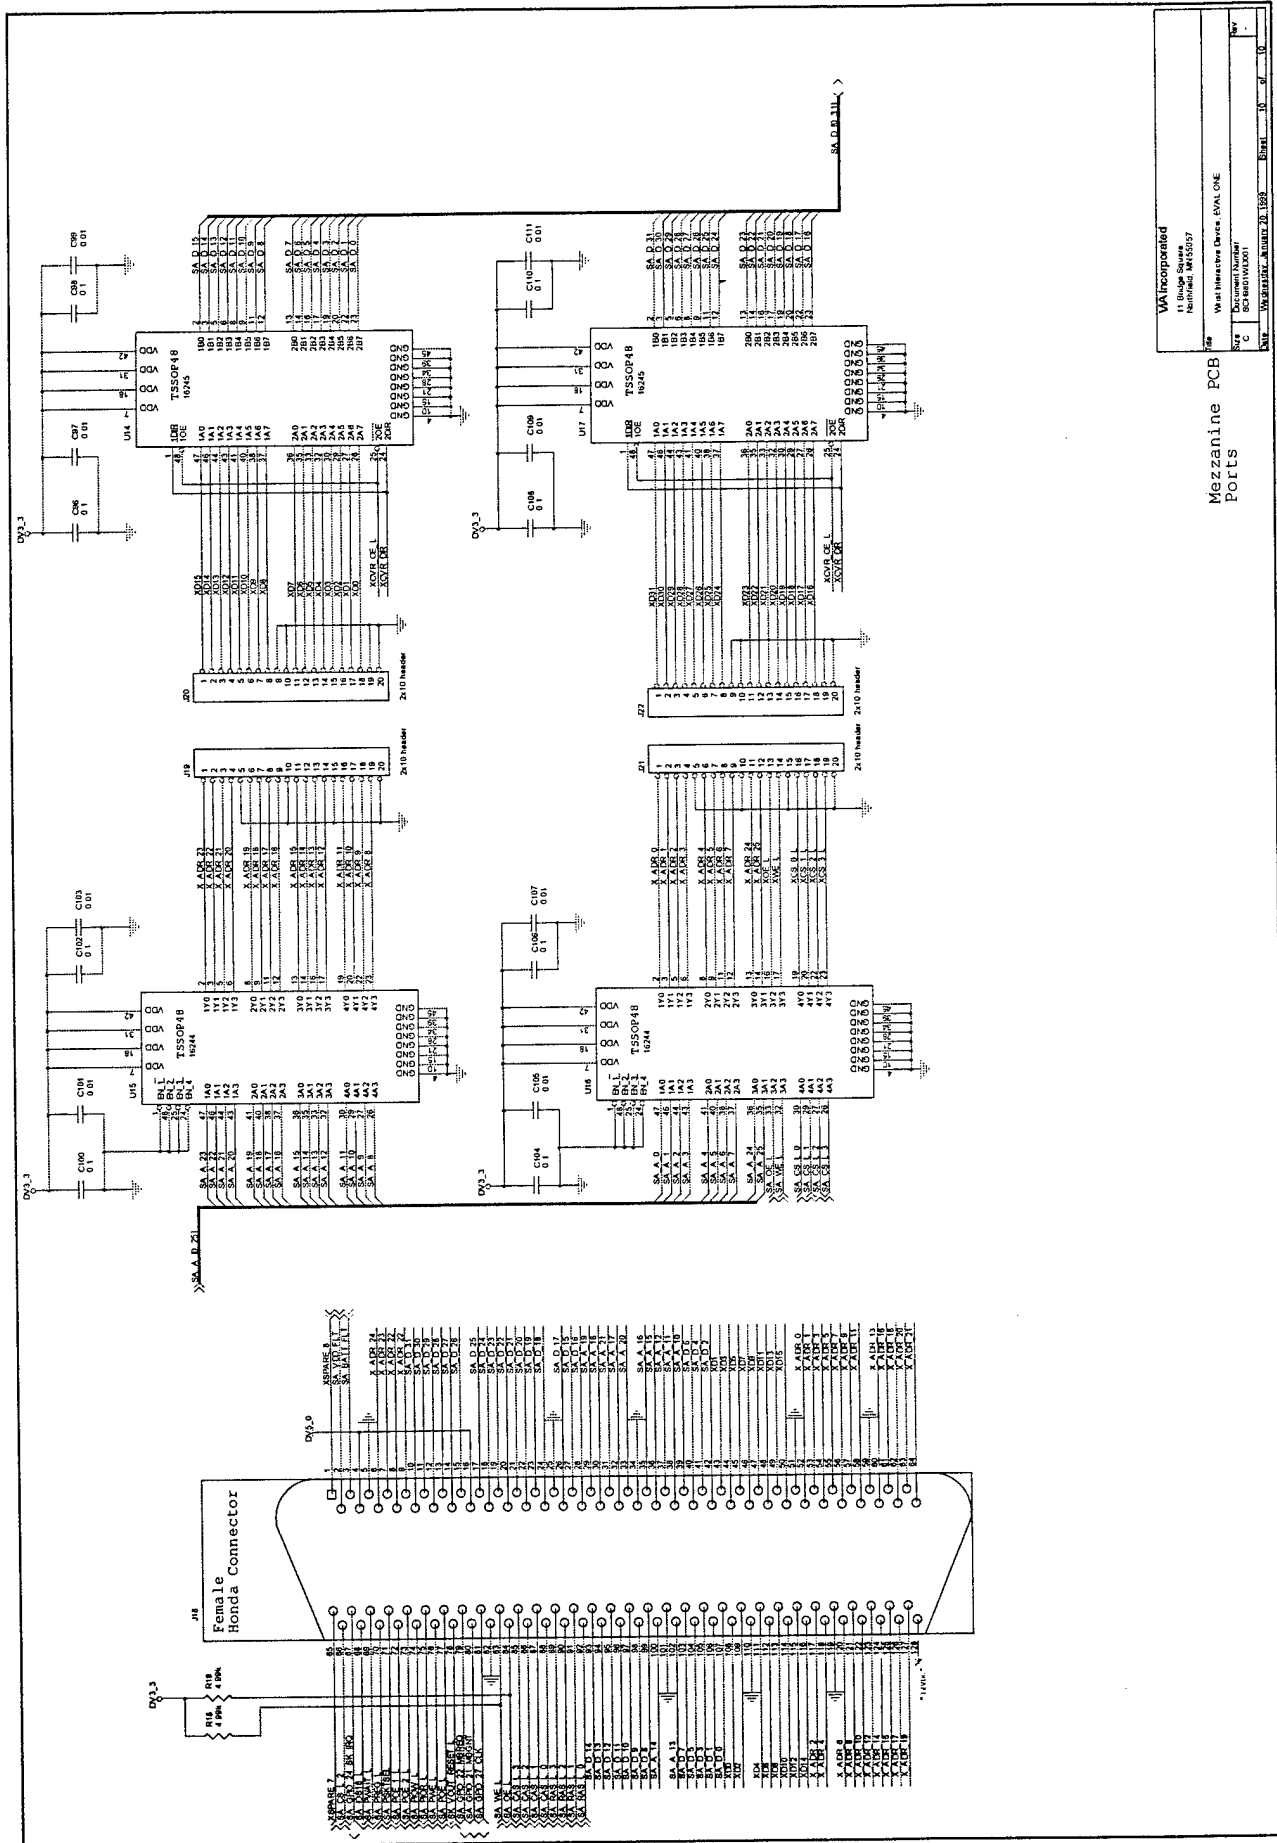


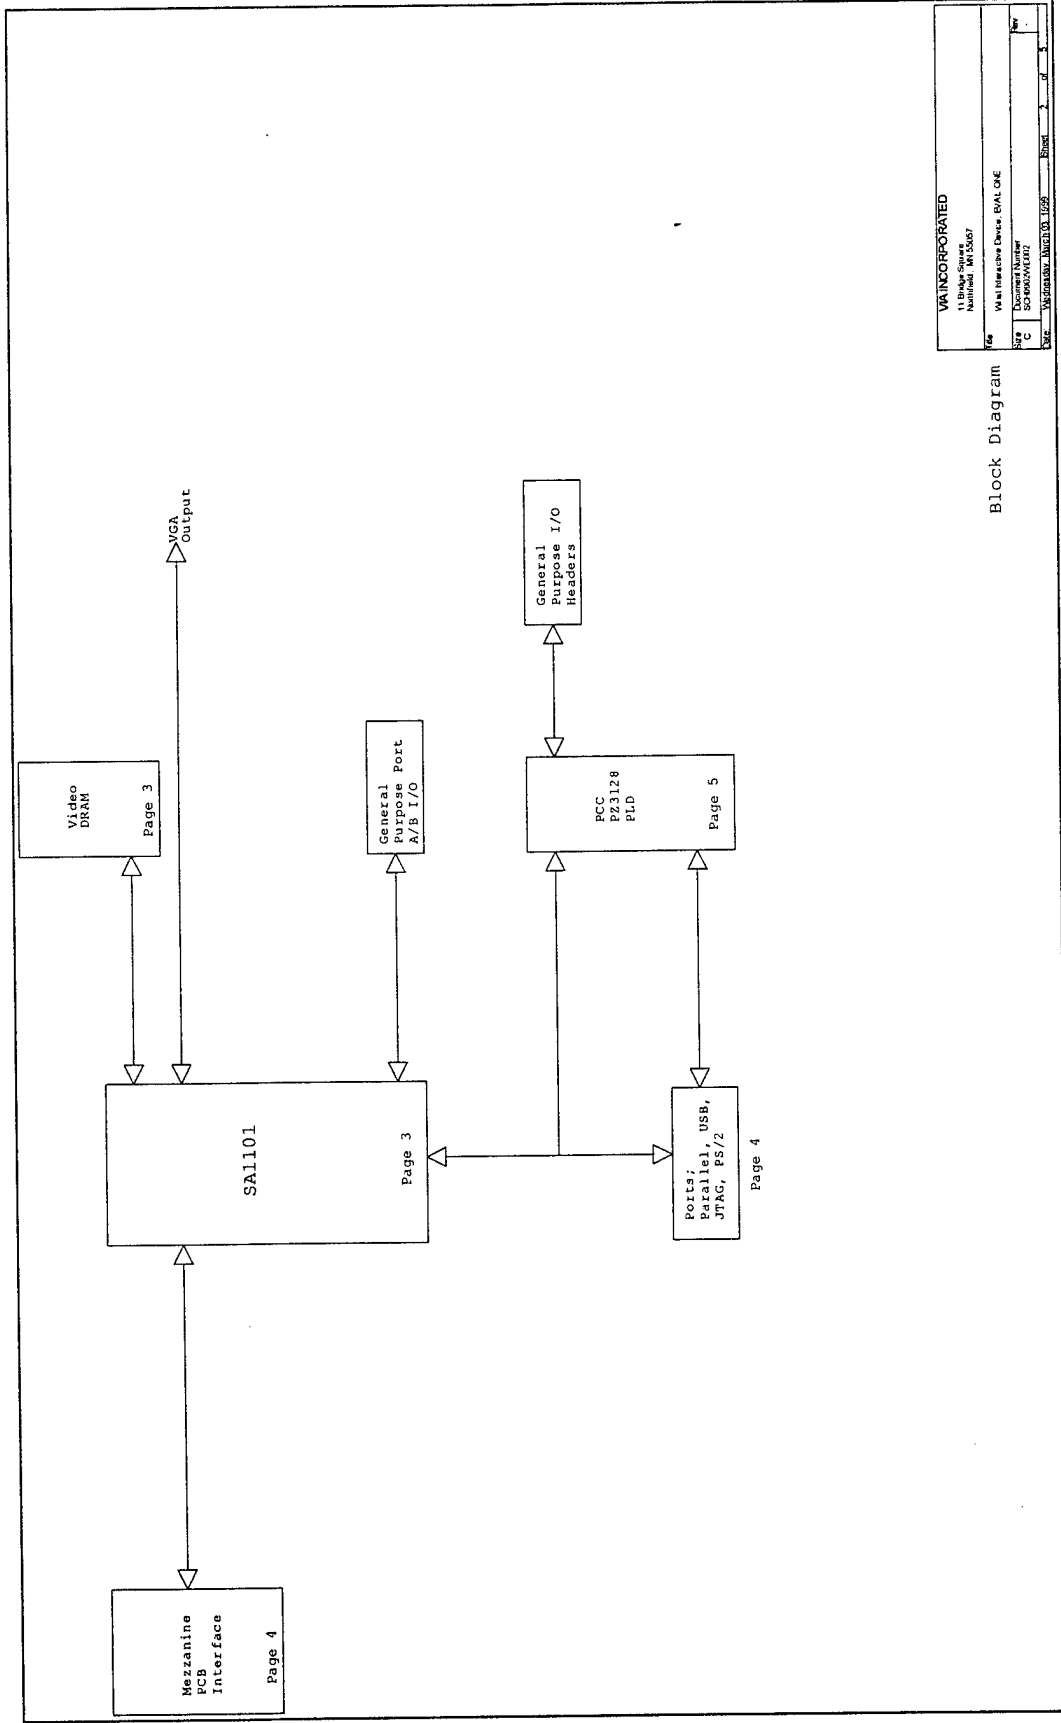


VIDEO
VA INCORPORATED
 11 Bridge Square
 Northfield, NY 55057
 Title: Wist Interactive Device EVAL ONE
 Size: B
 Document Number: SCH980WID001
 Date: Thursday, January 21, 1999
 Sheet: 7 of 10



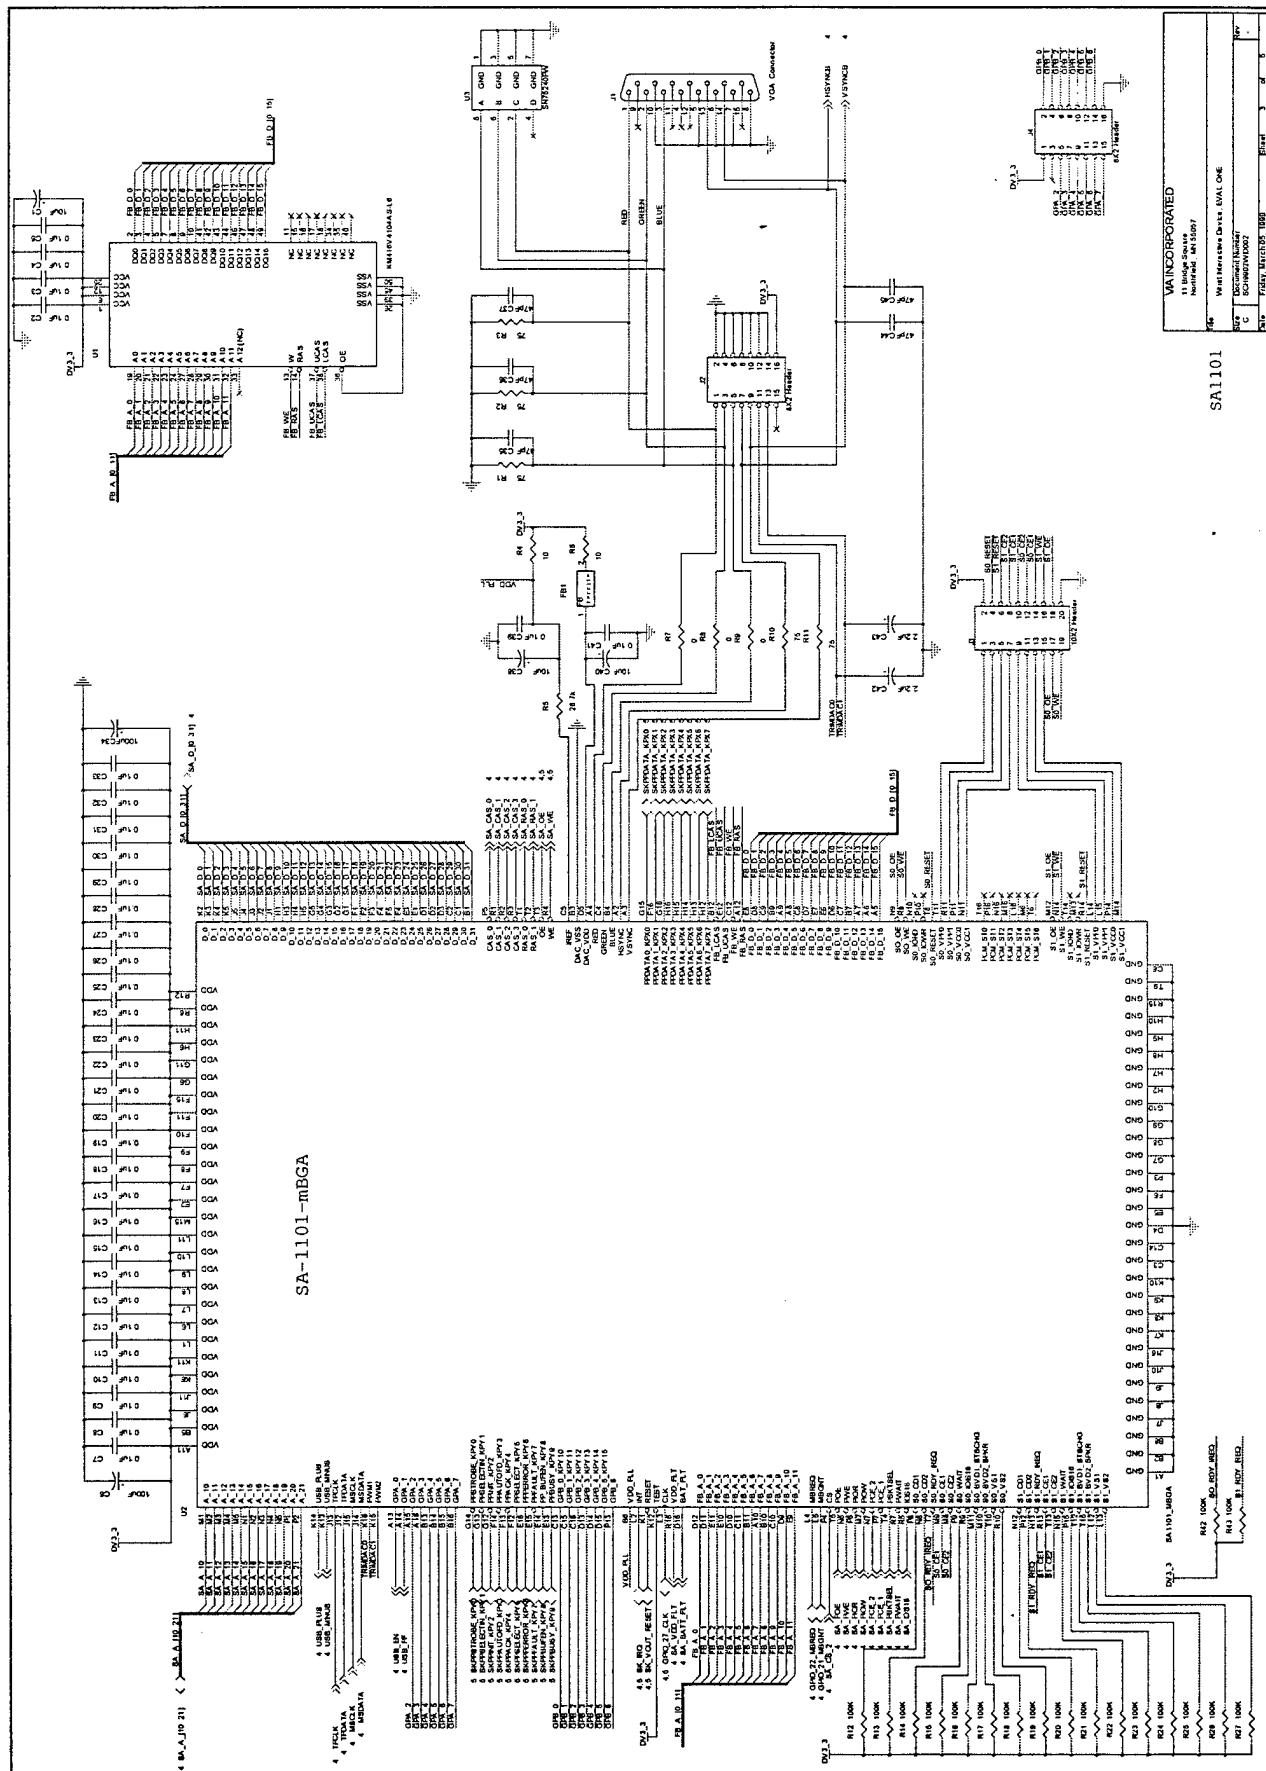
VA Incorporated	
11 Bridge Square	
Northfield, MN 55057	
File: Wisth Interactive Device, EVAL ONE	
Size	Document Number
0	SC98001WD001
Date	Wednesday, January 20, 1999
Sheet	9 of 10
Rev	-






VAT INCORPORATED	
11 Bridge Square Northfield, MN 55057	
File	Wat I/O Interface Device, BVAL ONE
Rev	1.0
Doc No	DOC 000001/01/02
Rev	1.0
Doc No	DOC 000001/01/02
Rev	1.0
Doc No	DOC 000001/01/02

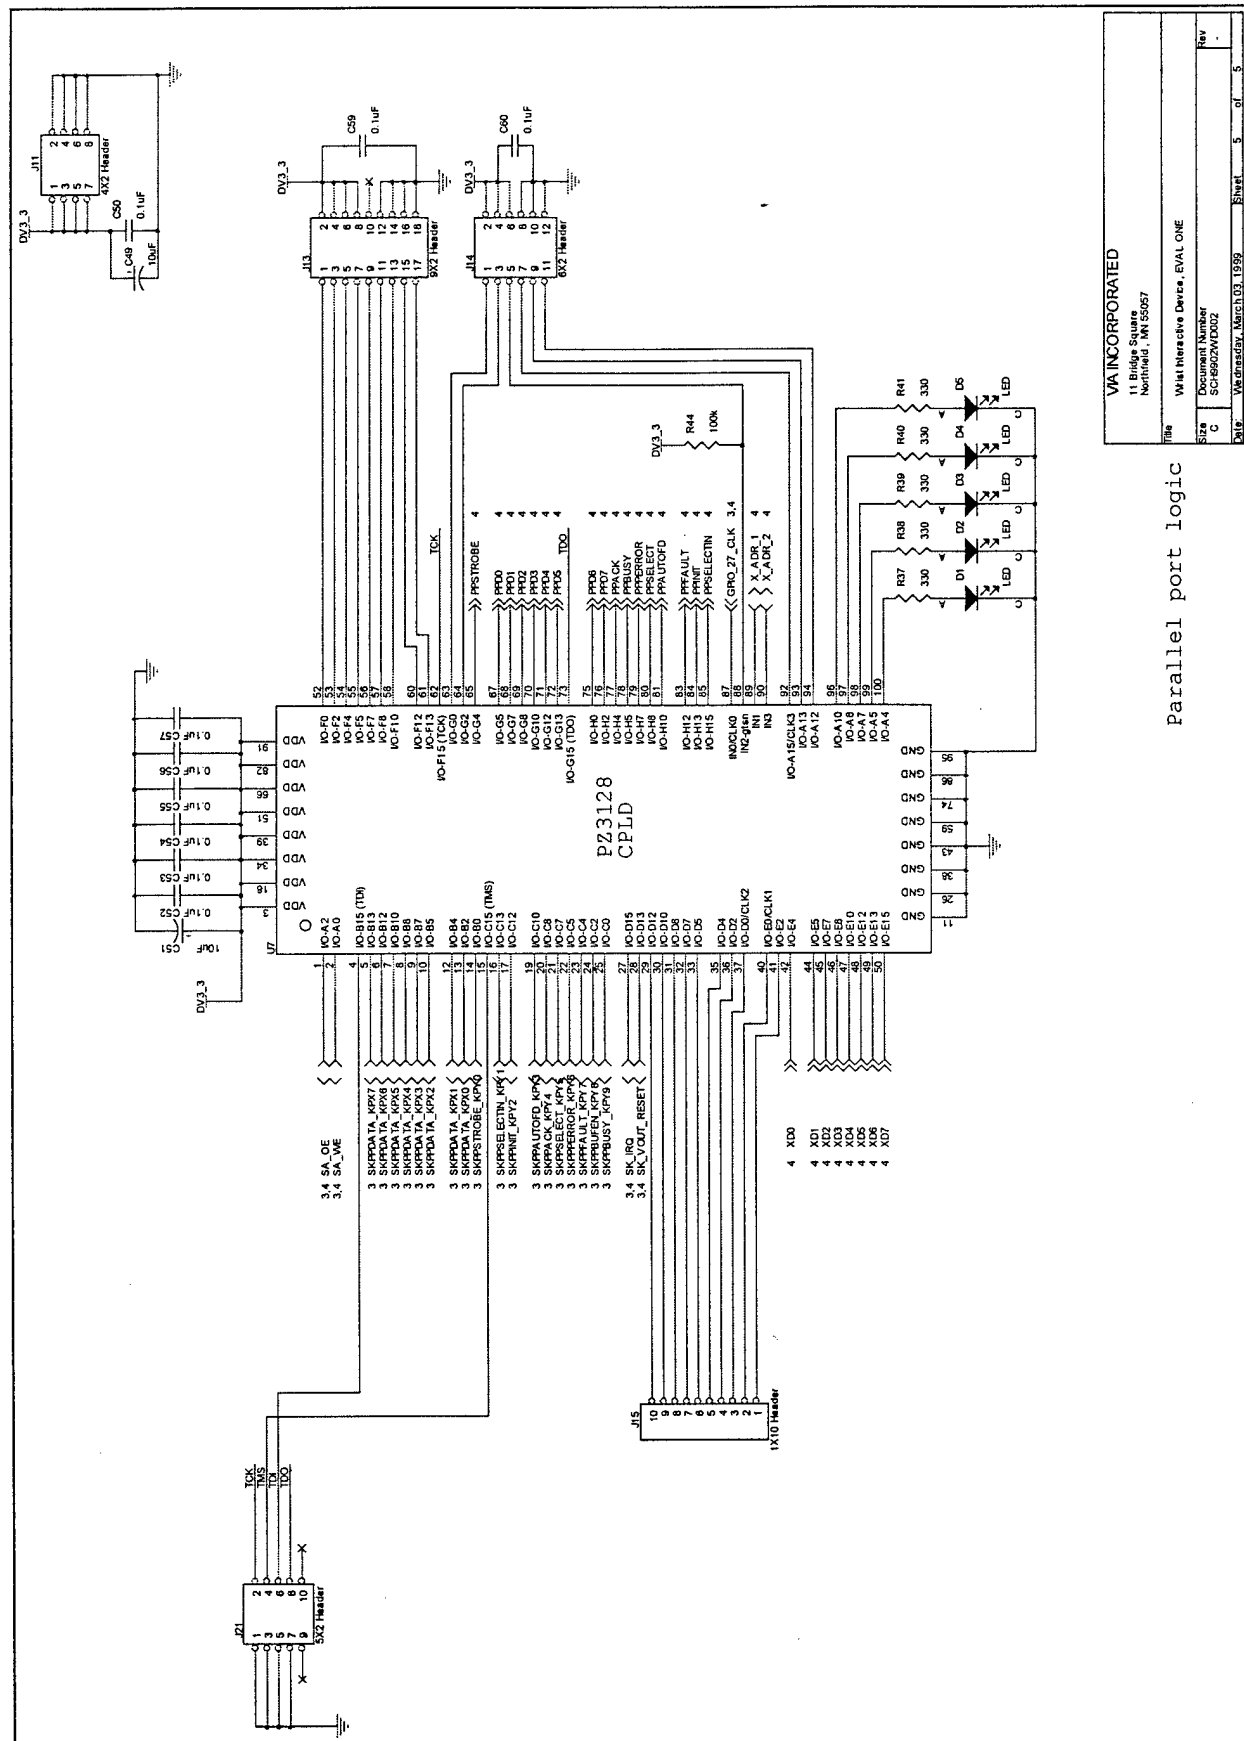
Block Diagram





Mezzanine Ports

VA INCORPORATED	
11 Broad Street, Suite 100	
Boston, MA 02107	
Tel: 617.552.1500	
Fax: 617.552.1501	
Email: sales@va-inc.com	
Web: www.va-inc.com	
Document Number: 0000000000	
Rev: 1.0	
Date: 08/28/2000	
By: 0000000000	
Check: 0000000000	
Part: 0000000000	
Rev: 0000000000	



Chapter IV

MAGNIFICATION OPTICS

Background and Preliminary Study

The magnification system for the WID poses some significant challenges. Depending on the type of display used the magnification and illumination optics will be entirely different. From the point of view of the magnification optics, all miniature displays ViA has examined fall into two broad categories: emissive or transmissive, and reflective. In the present report the emissive and/or transmissive type of display will be emphasized. However, in order to introduce ideas and concepts common to both, the Siliscape reflective display will be discussed in detail. Together with a general discussion of the main variables and concepts, this section proposes a preliminary design for the magnification system of the FED display.

ViA has considered two candidate lens forms for the monochrome green organic LED, 1280 x 1024, emissive microdisplay made by FED Corporation: a refractive aspheric singlet, and a spherical mirror lens. These have been analyzed in three configurations:

- 1) The singlet with off-axis, folded ray bundle
- 2) The same refractive singlet in a collapsible lens housing
- 3) A two-mirror reflective eyepiece, folded on-axis

Definitions and Technical Background

Independent Performance Variables

- a) Brightness. The display should be readable in daylight, though not in direct sunlight. This suggests a minimum visual brightness on the order of 300 nits (cd/m^2 , or 88 ftL). This is really bright and power consumptive. Therefore, for a longer battery life the system should have a 2- or 3-level brightness adjustment: 25% (normal indoors), 50% (bright room), and 100% (outside).
- b) Field of View. The field of view angle characterizes the magnification of the image relative to the object. (Since the object is placed at one focal length away from the lens, the magnification defined as the ratio of image location to object location is undefined as the (virtual) image is at infinity.) A 30° diagonal field of view is the equivalent of a 6.4" diagonal screen held at 12", a 5" x 7" screen held at 16", or a 15-inch desk monitor viewed from two feet. This is the smallest practical solid angle for a VGA-class screen for reading a page of 12-point text.
- c) Eye Relief. Distance from the lens plane at which the object can be viewed comfortably by the eye. The optics aperture increases proportionally, and the lens design becomes quite challenging, as the eye relief grows longer than the focal length. At 30 mm the nose and eyes clear the wrist well.
- d) Vignetting. Partial obstruction of the object being viewed through an optical system, usually around the edges, often caused by the optical system itself. Decreases the clear aperture.
- e) Exit Pupil. Diameter of the region perpendicular to the optical axis through which the entire object can be viewed. A 10 mm exit pupil for full-screen viewing at 50% vignetting should

be used to determine the lens's clear aperture. Of course the magnifier eyepieces do not have a real exit pupil, but the "eye box" for a complete diagonal will be limited by the lens's clear aperture, and/or the off-axis control of lens aberrations, and/or the light distribution.

Dependent Performance Variables

- a) Focal Length. Distance from the plane of the lens to the point on the lens axis where parallel incoming rays converge; inversely proportional to and a measure of the power of the lens. Real lenses have a complex geometry characterized by more than one focal length (Back Focal Length and Front Focal Length). In comparing optical systems it is common to utilize the average between the BFL and the FFL, called the Effective Focal Length (EFL), while still referring to it as simply the focal length. Fig. 4.1 shows the required focal length as a function of the active diagonal size of the object and for several desired diagonal field-of-view angles (= magnifications).
- b) Clear Aperture. Effective useable lens diameter. The geometry for the conditions above requires a 26 mm diagonal clear aperture. This is easy for the 37 mm EFL but more challenging for the 20 mm EFL.
- c) Power Consumption. The brighter the display the higher the power consumption. Although the power consumption is influenced by many display components, the power required by the LED illuminators for reflective displays is determined by the desired brightness at a rate of approximately 100 mW/100 nits. The power consumption of emissive displays can be similarly adjusted but the exact dependence is not yet known at the time of this writing.




Fig. 4.1 Required Focal Length as a Function of Object Diagonal and Desired Field-of-View Angle

Definition Sketch

Fig. 4.2 shows the main variables for the design of the 37mm EFL refractive aspheric singlet, for an emissive/transmissive display. This preliminary lens design has less than 2% distortion (pin-cushion), and worst-case of 1 diopter accommodation between tangential and saggital focus, in the corners. The index of refraction of polystyrene was used.




Fig. 4.2 Definition Sketch for Emissive/Transmissive Display

Singlet With Off-Axis Folded Ray Bundle

Ways of folding the optical path to allow most of the volume to be used at least twice have been examined. The best folding seems to be as shown in Fig. 4.3 below. The system is folded along the vertical direction, so in the picture the narrowest side of the screen and lens is shown. To achieve this folding we have shaved the lens's upper and lower edges until the pupil is approximately 40% vignetted. This liberty is taken because of the unique flexibility of the eye-wrist accommodation of position: users will position their wrist (and thus the exit-pupil on the eye) however needed to view the section of interest on the screen.

This approach has a major advantage: the exit axis is tilted some 40° from "vertical", allowing the user to adopt a much more natural wrist position for viewing. Both left- or right-hand mounting will have this benefit. An arbitrary volume for electronics and battery is indicated, along with an approximate wrist cross-section. The cross-section of the lens system is approximately 40 mm wide x 19 mm high, as shown. Most of the space inside the folded path is used twice, for a lens system volume on the order of 13.5 cm³ (assuming a depth out of the page of 2.5 cm).




Fig. 4.3 Singlet with Folded Axis

Same Refractive Singlet in a Collapsible Lens Housing

A second evaluation has looked at a variety of methods for collapsing the lens system for storage. Some were spindly assemblies of small girders; some were sealed bellows; most required some form of flexible storage bag to keep objects out of the viewer system; almost all were simply not on the path to a practical wrist display. The best approach is based on the field-proven collapsible drinking cup. Fig. 4.4 shows how the lens volume is collapsed along its axis.

There are five tapered and nesting rings to distribute the front focal length of the lens, allowing some 20% overlap. The lens is made round, to avoid any special orientation relative to the screen. It is deployed by pulling up and twisting to make a friction lock in the "focused" position. It is stored by shoving the lens down to a detented "click" to lock onto the base ring.

The lens-cell ring and the base rings are made of plastic or machined aluminum. The interfering rings between are stamped from stainless steel or some other tough metal, approximately 1 mm in thickness. The lens volume is reasonably sealed against penetration by water splash or dust in either the stored or deployed position. The structure is of course water-tight when deployed. Careful engineering of the detented rim in the cap should accomplish this for the storage seal also. The stored lens system will be very robust against direct blows and environmental assaults when stored. When deployed, the system can inadvertently be collapsed if hit against something

or leaned on. Prototype engineering design can explore detailed approaches to giving the “up” position a more detented and unbreakable interlocking.

One of the most difficult problems of a collapsing-lens approach is the swept volume: moist/hot/cold air is inhaled and exhaled by the system. In this preliminary sketch, no provision is made for cleaning the inside of the lens. In practice, the user must be able to remove easily the lens in the field and wipe the inside with a dry cloth to remove condensation.

In this configuration, the axis of the viewer is perpendicular to the wrist; the user must twist and lift his/her wrist near to its comfortable limits of movement. The height of the stored lens system is approximately 10.9 mm, with the lens vertex a bit above that. The base diameter of the (round) lens system is approximately 42 mm. The volume of the stored lens is approximately 15 cm³ and the volume of the deployed lens system is 26 cm³.




Fig. 4.4 Axially Collapsing Lens Storage (“Picnic Cup”)

Two-Mirror Reflective Eyepiece, Folded On Axis

For modest fields of view, a spherical mirror operated on axis can provide very good imagery, as a telescope objective or as an eyepiece. Fig. 4.5 shows the basic geometry, which is invented over and over with small variations. A spherical mirror of radius R has a focal length of $R/2$. For the WID the EFL is 37 mm, so the radius of curvature of the mirror is 74 mm. An eye positioned at the center of curvature as shown enjoys an almost aberration-free view of the objects lying on the dotted surface: zero distortion, zero coma or astigmatism, and zero vignetting of the field. However, two problems arise: the object (e.g., and LCD) is generally flat, and it gets in the way of the view.




Fig. 4.5. Basic Imaging—Spherical Mirror

The flat object can be in focus only for a circular zone in the field of view (FOV), and the curvature of the spherical field limits the useful FOV to approximately 40° for a young viewer, 30° for the middle-aged, and less than 20° for old eyes. By inspection, one can see that the radius of curvature is equal to the focal length: half the mirror radius. Undeterred by this, new entrants to the head-mounted display and viewer fields continue to proffer breakthroughs in fields up to 60° or more, for which the field curvature is hopelessly too large to accommodate.

One of the best early applications of this approach was developed at S-TRON Inc. for Air Force and Army helmet-mounted displays, as shown in Fig. 4.6. The object is removed from the field of gaze (here shown with a little vignetting at the edge of the FOV), at the cost of introducing a 50:50 beam splitter. This was used to good advantage in the Soldier's Integrated Protective Ensemble (SIPE) developed by S-TRON Inc. and was later directly adopted by Virtual I/O in its widely sold "I-Glasses" HMD product line.

In Fig. 4.6 the field curvature is obvious for this modest 30° FOV, requiring accommodation of the order of 1 diopter from center to corner. The image light bounces once and is transmitted

once through the beam splitter. If the source is monochromatic then a dielectric beam splitter can approach the theoretical maximum of $50\% \times 50\% = 25\%$ system transmission. If the source is broad spectrum, a simple metal beam-splitter coating will reflect and transmit approximately 40%, for a practical transmission of 16%, source-to-eye.




Fig. 4.6 Folded, Still On-Axis (as in Edwards's SIPE, Virtual I/O HMDs)

The system is obviously too tall as deployed, for application to the WID display. However, folding up this system appears to be unproductive. While all the parts can certainly be hinged and repositioned, the large footprint when folded and the necessary dust-bag enclosure make it unattractive to pursue for the wrist display. Under contract to Kaiser Industries, Arthur Berman invented an interesting variation of Fig. 4.6, in which the optical axis was folded back on itself. This on-axis fold is shown in Fig. 4.7. The same curved focal surface is shown on the object.




Fig. 4.7 Kaiser On-Axis Fold (Berman)

By using a cholesteric liquid crystal mirror on the flat shown, Berman achieves 100% reflectance of circularly polarized light, which reverses its sign on reflection from the beam splitter, and is completely transmitted through the cholesteric layer. For the narrow wavelength for which the cholesteric liquid crystal mirror is tuned, the system transmission is decreased only by the transmission/reflection of the beam splitter. If the source light is unpolarized the system undergoes an additional halving of energy for circular polarization. In the case of an LCD source the exit light is already linearly polarized, and convertible to circular polarization without losses.

For the unpolarized light of the organic LED, the practically achievable green transmission will be approximately $(40\%) e^3 = 6.3\%$. The light traverses the interior cavity three times. The minimum axial length of the triply-folded system is $1/3$ of the back focal length, or approximately 12 mm.

Siliscap Corporation owns a series of patents using essentially the same design as the original "pancake window", but with the object imbedded in the first surface. One of their early patents is sketched in Fig. 4.8. Here we see that the interior space between mirrors is traversed five times, making the minimum axial length near to $1/5$ of the focal length ($37/5 = 7.4$ mm). For a reflective object source this system has the special requirement that the object diagonal be smaller than the eye's pupil (preferably less than 3 mm diameter) to permit the system to operate without crippling vignetting blocking operation near the optic axis. This of course places an extreme requirement on the power required to have the tiny object subtend a 30° field of view, as indicated below.




Fig. 4.8 Original Siliscap System Concept

The light makes two reflections and one transmission at the (metal) beamsplitter, so the system transmission is practically $(40\%) e^3 = 6.3\%$.

The Siliscope eyepiece, which has been brought to market, is of a completely different form. Recognizing that the object should be removed from the interior space of the lens system, Siliscope has essentially returned to the original beam splitter design of Fig. 4.6 above. As shown in Fig. 4.9, the beam splitter is a complex prism with total internal reflection over a small range of angles for light leaving the object. As can be seen by inspection, the illuminating source cannot easily be specularly reflected into this narrow acceptance cone, so a completely different kind of liquid crystal display is required: the LCD either scatters the incident light (for white) or specularly reflects and dumps the incident light (for black). This approach will work well with self-radiant sources, such as the organic LED displays.




Fig. 4.9 Siliscope System as Marketed

Advantages of the Siliscope Lens:

- Effectively 100% transmission, making the display very bright along an individual ray.

Disadvantages of the Siliscape Lens:

- Longer along its axis than the form of Fig. 4.7: the interior length of the system is approximately 70% of the picture height, vs. 1 x the picture height in Fig. 4.6. In the case of the FED 1280 x 1024 display (15.4 x 12.3 mm), the interior length of this lens form would be 8.6 mm, to which would be added the thickness of the reflecting/refracting components.
- Less magnification: the optical path length through the prism cluster is approximately twice the focal length for the similar systems above drawn on the same scale and thus the power of the lens is halved.
- Less light gathering: the solid angle of light gathered into the pupil includes approximately 10% - 20% of the light gathered by the foregoing examples. Thus the image brightness of the Siliscape lens is indistinguishable from that of the other designs considered.
- Heavier: interior of lens system in solid glass.
- Bulkier side extension (of glass prism) for image source.
- Conventional LCDs will have a very poor performance.

Finally, Fig. 4.10 shows a new concept for an emissive display developed by DisplayWear to achieve some of the desiderata of compactness, rigidity, ruggedness, and low cost. Here two concave half-transmissive mirrors sum to a power of 27 diopters, or a focal length of 27 mm. The minimum thickness for this triply traversed lens is $EFL/3 \approx 12$ mm, allowing 1 mm of focal length outside for mirror thickness. As shown in the second part of the figure, this lens system could be collapsed to a thickness of less than 6 mm if that were vital, using the tapered-ring approach of Figure 4.4 above.

Summary

Table 4.1 compares the relevant performance variables for the different lens systems examined.

Table 4.1 Summary of Findings							
Performance of Lens in Fig:	4.3	4.4	4.6	4.7	4.8	4.9	4.10
Focal length (mm)	37	37	37	37	37	74	37
Field of view (deg.)	30	30	30	30	30	15	30
Eye relief (mm)	35	35	35	35	35	35	35
Vignetting (Top & Bot)	Yes	No	Yes	No	No	No	No
Green transmission (%)	100	100	25	12.5	12.5	100	12.5
Avg. white transmission (%)	100	100	16	6.3	6.3	100	2.6
Height, deployed (mm)	18.8	40	30	13	8.2	11-15	12
Height, collapsed (mm)	-	11	-	-	-	-	6
Relative brightness*	1	1	0.16	0.06	0.06	0.15	0.06

*using unpolarized image source

The system of Fig. 4.3 is more ergonomic, while that of Fig. 4.9 is made of solid glass and is therefore relatively heavy. The systems of Figs. 4.3 and 4.4 are of particular interest for their compactness and high image brightness. The system of Fig. 4.10 is distinguished by its low

profile when collapsed, and by its mediocre transmission. The other systems do not recommend themselves to use in a practical wrist display.




Fig. 4.10 DisplayWear's Biconvex Mirror Lens

Design of Optics Mock-Up for Emissive Display

Lens

An aspheric biconvex plastic design was originally designed by DisplayWear and delivered satisfactorily. After examining the performance, it became apparent that the singlet lens showed about 8% pincushion distortion. While this would normally be satisfactory, DisplayWear elected to use a lens they had previously designed for use in a head-mounted display. This lens is a diffractive aspheric element, with over \$35,000 in the tooling from which it was made. It was fortunate to have this available, in that it provides distortion-free, color-corrected, high-resolution imagery. This lens is a two-piece lens, with a field flattener at the focal plane. The overall focal length of the combination lens is approximately 36 mm. Thus, the 19.7 mm

diagonal of the image source in the delivered display subtends 30.6° . DisplayWear has implemented this two-element lens in the WID mock-up, at an increase of some 5 mm thickness, because in their opinion it gives a much improved demonstration of the “look and feel” of a wrist-mounted display over the singlet originally designed. Fig. 4.11 shows the lenses and the WID mock-up in extended and collapsed conditions.




Fig. 4.11 Collapsible Magnification System for Emissive Wrist-Mounted Display

Ring Fabrication

The loose conical rings are made by spinning an aluminum cylinder on a steel mandrel, and cutting sections at calculated axial distances. The aluminum alloy is work-hardening so it is quite stiff in the final 0.021" thickness.

Assembly

This optomechanical demonstration prototype does not use the actual OLED microdisplay. The OLED microdisplay is simulated by a transparency image of a computer display viewed against a green background.

The uppermost ring element is machined, not spun from sheet like the other rings. Assembly requires to place the concentric rings in order and to attach the telescoping base ring to the display base. The now-captivated rings terminate in the uppermost ring element. The last element to be added is a cap ring which captivates and seals the lens.

Backlighting

The backlight is white: two filament lamps flood the diffusive reflector on the green base floor and transilluminate the screen object. Two axial lead 5 volt bulbs are used. Running them at 6 volts, the current is 82 mA per bulb, and the life is reduced from 25,000 hours to 3000 hours. Thus the total draw is about 160 mA (1 W total). It is important to note that the light sources in this concept model will not be needed by the OLED microdisplay since it is emissive.

Batteries

Initially the working mock-up was designed to use internal batteries. The stored energy should have given approximately 20 minutes of use before changing batteries. DisplayWear discovered late in their development that the wafer batteries had a current limit well below what was needed. To meet the schedule they substituted the original, internal coin cell for an external battery module, and used a relatively large lithium 6 volt battery which is used to power 35 mm camera motors (type 23-178).

Although a smaller battery might have been chosen for the external pack, DisplayWear went ahead with this high-energy pack for reasons of schedule and predictive efficacy. Changing the battery is done by removing the screws which hold the battery case cap. If the battery is hard to remove, a paper clip can be used to push the battery out through a hole in the top of the case.

It is expected that two thinner lithium-polymer or lithium-vanadium pentaoxide prismatic batteries (to be discussed in Chapter V) will be used in the final prototype. These cells will be shaped to conform to the wrist band and will not be as bulky as this one.

Color Scheme

All parts are anodized flat black. The telescoping rings are hard-anodized for scratch-proof, low-wear usage. A matching black wrist strap is installed.

Display Screen Image

After receiving a screen print from ViA, DisplayWear scaled this monochrome image on 35 mm film to act as the apparent screen.

On-Off Switch

A momentary off-on switch has been provided, as it tends to save battery power. In the final model the switch may toggle from ON to OFF.

Operation

The display is opened by twisting and lifting the top ring, which holds the lens. It is lifted until the rings engage, then twisted slightly to seat and seal the rings. Similarly, the display is closed by twisting the top cap while collapsing the display, and twisting the cap when it bottoms on the lowest ring, to seat and seal the display in the closed position.

It is important not to force the rings in tension as they can be pulled past each other. In the event that someone pulls the rings apart, do not try simply to force them back as damage to the interfering lips will result. Take the stack apart, and reassemble in reverse order.

Design of Optics Mock-Up for Reflective Display

DisplayWear has succeeded in designing a folding optics train for a LCOS, reflective microdisplay, as shown in Fig. 4.12. The design parameters were chosen as follows:

Table 4.2 Design Parameters for Reflective Display	
Brightness	approximately 30 ftL (100 cd/m^2) in the demonstration prototype, approximately 80 ftL (270 cd/m^2) in the production system (depending on power available).
Color	The color depth of the display will depend on the host processor and on the display chip itself.
Field of View	Eye relief scales inversely with field of view and object size. Because the LCOS display generator is 60% of the diagonal size of the OLED pixel generator considered earlier, the field of view is decreased to a small, but still useful, 25° diagonal.
Focal Length	Thus, for the 11.9 mm diagonal object, the 25° field of view will require a lens focal length of 26.8 mm.
Eye Relief	A traditionally designed eyepiece is telecentric: the chief rays between the object and the lens are parallel to the axis. For this case the eye relief equals the focal length. In the present case a 30 mm eye relief can be achieved by careful placement of the light source to obtain a divergence in the chief rays from the image points.
Exit pupil	10 mm
Lens dimensions	27.5 mm lens diameter, with a 25 mm clear aperture. The lens will be designed to have one aspheric surface to minimize distortion and to balance astigmatism.
Power	See Table 2.1, assuming the CMD display.

As noted above the lens is illuminated to operate non-telecentrically, to extend the eye relief from its focal length (27 mm) back to a more useful 38 mm. The illuminator is designed to approximate the function of an integrating sphere, sending light from its aperture in a Lambertian exit pattern. The light is focused forward via a Fresnel lens which focuses the light approximately on the beam splitter. The light reflects and diverges downward to reflect from the microdisplay surface along the directions of the chief rays from each point.

The exit port of the illuminator is closed by a film pre-polarizer oriented for S-polarization relative to the polarizing beam splitter. The eye lens is preceded by a P-polarization post-analyzer element on a glass plate (the black line below the lens), which cleans up the blue and any residual leaking S-polarized light.

The beam splitter is broadband and polarizing in a plastic sheet format, oriented to reflect S-polarization. It is sufficiently flat to *accurately* transfer the illumination from the LED source integrator. The requirement for flatness in reflection is considerably relaxed by focusing the source on the surface of the beam splitter, in order to achieve the source divergence.




Fig. 4.12 Collapsible Magnification System for Reflective Wrist-Mounted Display

The beam splitter is held in a plastic frame. The plastic frame terminates at the top in two ears formed as cam followers. These ears slide along the glass of the upper analyzer, as it presses down the folding components. These ears are Teflon-coated to avoid scratching the glass. A shaft through the foot of the frame is captured in a bearing set (not shown) and a spiral spring (not shown) presses the frame and (through linkage) the integrated Fresnel lens/illuminator to rotate in a counter-clockwise direction. A stop at the foot of the beam splitter frame keeps the deployed beam splitter at its preferred 45° from the lens optic axis, and the lens/illuminator in its erect position on the reflected optic axis.

When the lens is lowered the top of the beam splitter frame follows the flat surface of the filter, folding both the beam splitter and lens-illuminator over to a minimum-height stored position. It was especially pleasing to find a way to work this out, since in earlier work such compact storage of the large components appeared unlikely. Considerable space is left for an electronics bay, amounting to approximately 4 cm^3 , not including space above the plane of the display element.

When the system is closed it is sealed watertight. When it is opened it is sealed watertight. In between it is open to the atmosphere, to prevent buildup of pressure or vacuum. It locks open or closed by friction, needing no complex detents or latches.

This bright, full-color, 25° FOV video display will satisfy the visual needs of the wearable computer user, and the environmental and human factors as well. Its unique folding and collapsing mechanics will allow it to be worn as an ordinary wrist unit, and deployed for instant use by a gentle twisting.

Lens

At midterm DisplayWear reported that the lens, an aspheric biconvex plastic design, had been delivered twice with different disabilities which necessitated its return for rework. The lens was subsequently returned and found marginally acceptable. While it appeared to have been made with reasonable care, the corners were too fuzzy.

The central problem with designing a minimum-height lens is that at this power and field of view the lens should really be split into a doublet. DisplayWear attempted to do it in a singlet to save the vertical height; however, in DisplayWear's opinion the attempt was not satisfactory. Therefore, as for the emissive display, DisplayWear converted the system design to utilize the same excellent lens adapted from an earlier program. The focal length of this lens is somewhat longer than the intended lens: 31 mm vs. 27 mm. Thus, the 12 mm diagonal of the image source in the delivered display subtends 22° , something less than the target of 25° .

Ring Fabrication

The same collapsible "picnic-cup" design of the emissive display is used here. Loose conical rings are made by spinning an aluminum cylinder on a steel mandrel, and cutting sections at calculated axial distances. The aluminum alloy is work-hardening so is quite stiff in the final 0.031" thickness. The same mandrel was used for fabricating rings for both collapsible magnification systems. As initially delivered the rings were found to have too much interference. To meet delivery schedule DisplayWear simply cut the ring wall thickness to

approximately 0.018", and got them hard-anodized with enough surface build-up to provide 0.100" overlap at each ring joint.

Hand-making the first one of these tapered-ring telescoping stacks is an unnerving fabrication task, and requires extensive hand-fitting and re-work. For illustration, at the taper angle of 1.5° (3° apex angle), the hard-anodized thickness of 0.0013" corresponds to a relative movement of the tubes of 0.1" axially. The anodizer was quite taken aback by a thickness tolerance of 0.0013" to 0.0015" allowable. Making the rings in production would be much simpler and quite reproducible, with each step following a formula worked out in manufacturing engineering.

Assembly

It will be noticed that the assembly of the current physical model is somewhat taller than defined (and manufactured) at midterm by the addition of a 6.5 mm spacer, to accommodate the longer focal length of the lens utilized. Fig. 4.12 reflects the initial conception and the intended design, which will be demonstrated with the next physical model.

This optomechanical demonstration prototype does not use the actual light source and optics of the folded reflective microdisplay. The mechanical parts of the reflective microdisplay are however included and shown to deploy when the eye lens is erected, and they fold up for storage. A coil spring (not shown) at the base of the beam splitter frame erects the beam splitter when the eye lens is raised. It is mechanically stopped at an inclination of 45° . Similarly, a coil spring (not shown) at the base of the combined collimating lens and integrating-sphere mock-up of the LED frame-sequential tricolor source erects that subassembly when the eye lens is raised. It is mechanically stopped at an inclination of 90° , or vertical. The frame-sequential illuminator shape is taken from the existing source in the Colorado Microdisplay (CMD) reference system. In the actual system it contains three color LEDs and acts as an integrating sphere source to mix the light. The Fresnel lens transfers this light into the system, focusing the light source approximately on the beam splitter.

A smooth rounded surface at the top of the beam splitter frame acts as a cam follower, running on the lower surface of the eye lens. A smooth rounded surface at the top of the collimator/light-source subassembly also acts as a cam follower, running on the lower surface of the beam splitter. The strength of the springs has been adjusted to be very light, and no visible wear is anticipated to occur on either optical surface acting as a cam. For ease in viewing the interior parts of this optomechanical prototype, the beam splitter element was left out. This permits the interior parts and their functioning to be observed as the lens is lowered to storage.

The uppermost ring element is machined, not spun from sheet like the other rings. Assembly will be by placing the concentric rings in order, and attaching the telescoping base ring to the display base. The now-captivated rings terminate in the uppermost ring element. The last element to be added is a cap ring which captivates and seals the lens.

Backlighting

The back light is white: two filament lamps flood the diffusive reflector on the base floor and transilluminate the screen object. This has been one of the more challenging parts of this detailed design, since the choices of physical size of both light bulbs and batteries are quite limited.

Two axial-lead 5 volt bulbs are used. Running them at 6 volts, the current is 82 mA per bulb, and the life is reduced from 25,000 hours to 3000 hours. Thus the total draw is about 160 mA (1 W total). It is important to note that the light sources in this concept model are some 40 times more power-consumptive than will be the triple LEDs used in a reflective microdisplay. This operating mock-up of the deployable wrist display uses a back-lighted film transparency instead of Kohler¹ illumination of the reflective microdisplay. In other words, this diffuse hemispheric back-lighting of the image transparency floods the transparency, instead of using the same path as the closely-controlled beam illumination of the reflective microdisplay.

Additionally, the image had to be bright. The high brightness is included in the engineering concept model to ensure that the user focuses on the important variables (size, folding, focus, image) and not on the (selectable) image brightness as a value determinant.

In manufacturing a working model four LED strips should be used, each consuming 30 mA. At maximum brightness (outdoors) all 4 would be used and the current would be 0.12 A. At lowest brightness (indoors) only one would be used, at 30 mA.

Batteries

Initially the working mock-up was designed to use internal batteries. The stored energy should have given approximately 20 minutes of use before changing batteries. DisplayWear discovered late in their development that the wafer batteries had a current limit well below what was needed. To meet the schedule they substituted the original, internal coin cell for an external battery module, and used a relatively large lithium 6 volt battery which is used to power 35 mm camera motors (type 23-178).

Although a smaller battery might have been chosen for the external pack, DisplayWear went ahead with this high-energy pack for reasons of schedule and predictive efficacy. Changing the battery is done by removing the screws which hold the battery case cap. If the battery is hard to remove, a paper clip can be used to push the battery out through a hole in the top of the case.

It is expected that two thinner lithium-polymer or lithium-vanadium pentaoxide prismatic batteries (to be discussed in Chapter V) will be used in the final prototype. These cells will be shaped to conform to the wrist band and will not be as bulky as this one.

Color Scheme

All parts are anodized flat black. The telescoping rings are hard anodized for scratch-proof low-wear usage. A matching black wrist strap is installed.

Display Screen Image

We have received two sets of candidate screen prints, for inclusion. These have been scaled on 35 mm film to act as the apparent screen.

¹ Kohler illumination is used in high-quality instrumentation such as microscopes and the best head-mounted displays. In this type of illumination, the "filament" (light source plane) is imaged into the pupil of the user. Thus, the exit pupil is filled with light and no light is wasted. In contrast, a diffuse backlight throws the light energy in all directions, including the location of the exit pupil.

On-Off Switch

A momentary off-on switch has been provided, as it tends to save battery power. In the final model the switch may toggle from ON to OFF.

Operation

The display is opened by twisting and lifting the top ring, which holds the lens. It is lifted until the rings engage, then twisted slightly to seat and seal the rings. Similarly, the display is closed by twisting the top cap while collapsing the display, and twisting the cap when it bottoms on the lowest ring, to seat and seal the display in the closed position.

It is important not to force the rings in tension as they can be pulled past each other. In the event that someone pulls the rings apart, do not try simply to force them back as damage to the interfering lips will result. Take the stack apart, and reassemble in reverse order.

Chapter V

BATTERY TECHNOLOGY

Introduction

To provide the expected average power the WID will rely on a rechargeable battery. The main goals of battery technology are to

- 1- Maximize energy density per unit volume
- 2- Maximize energy density per unit mass
- 3- Maximize number of recharge cycles
- 4- Optimize voltage discharge profile and current capacity (discharge rate)
- 5- Meet safety criteria

There are three major types of battery technology:

1. Consumer Batteries. Typically purchased at retail outlets for household applications, such as clocks, smoke detectors, radios, cameras, toys, flashlights, and cassette players. Single-use, disposable ("primary") alkaline and zinc-carbon batteries dominate this market. Rechargeable ("secondary") batteries are beginning to penetrate the consumer battery market. Most rechargeables are nickel-cadmium, but rechargeable alkaline batteries have also entered the market.
2. Automotive and Industrial Batteries. Automotive and industrial batteries are currently low-growth markets relying on lead-acid batteries, a mature technology. Substantial efforts are now underway to develop new battery technologies to power electric vehicles, such as fuel cells.
3. High-Performance Rechargeable Batteries. These are primarily used for portable products, such as camcorders, cordless power tools, household appliances, notebook computers and cellular phones. This is a fast-growing segment of the market due to the increasing number of "mobile" workers and consumer demand for portable electronic devices.

It is impossible to single out any one battery type as the "best", since each application has its own set of requirements. Fig. 5.1 shows how different battery technologies have prospered in different markets. It is clear that the WID will require a battery of the third type, which falls somewhere in the middle of the figure.

When comparing different high-performance rechargeable batteries for the WID, important characteristics beyond the main goals listed above include:

- Operational Temperature Range - The temperature range in which the battery can be discharged/used
- Cost - matters to just about all but the most specialized of users
- Other aspects:
 - Commercial Availability
 - Storage Energy Dissipation Rate (e.g., Ni-Cads lose appx. 1% of their energy level each day)
 - Form Factor (e.g., thin films are easier to integrate into a system design)




Fig. 5.1 Current Market for Batteries

Over the past twenty years battery technology has evolved toward higher energy densities. Fig. 5.2 compares several main battery types. The most striking feature of this chart is the great difference in energy density between lithium-vanadium pentaoxide and all the other types. This can be attributed in part to the fact that for the $\text{Li-V}_2\text{O}_5$ only laboratory figures are available, while the other types are actual production specs of commercialized products. Thus, while $\text{Li-V}_2\text{O}_5$ batteries promise to yield the highest energy densities, the numbers shown in this chart are somewhat optimistic.

Lithium-based batteries have the best energy density, the highest voltages, and high discharge rates. However, with these cells it is not always easy to meet safety requirements. ViA has examined three types of lithium batteries: lithium-ion, lithium-polymer, and lithium-vanadium pentaoxide ($\text{Li-V}_2\text{O}_5$). Lithium-ion and $\text{Li-V}_2\text{O}_5$ utilize a liquid electrolyte, whereas lithium-polymer batteries rely on a solid polymer to conduct the charges between the cathode and the anode. In battery terminology the convention is to place the anode (-) on the left and the cathode (+) on the right. Thus, in $\text{Li-V}_2\text{O}_5$, Li is the negative electrode or anode while V_2O_5 is the positive electrode or cathode. During discharge, lithium ions migrate from the negative lithium anode, through the electrolyte, to the positive cathode. During recharge, they migrate back.




Fig. 5.2 Energy Densities for Different Battery Types

Li-Ion Batteries

For many years the main obstacle to developing commercial rechargeable lithium cells was making the recharging safe. When lithium metal is used in conjunction with a liquid electrolyte, during recharge the lithium ions deposit on the anode's surface according to a crystallization process that favors the growth of irregular and tree-like structures that protrude from the metal surface. These crystals are called dendrites and have a very negative impact on battery performance and safety. Once crystallized, in fact, the lithium is not available any more for the discharge cycle and remains trapped in the dendrites. To reach the required number of recharge cycles, therefore, the anode must be supplied with considerably more lithium metal than would otherwise be necessary, decreasing the gravimetric energy density of the cell. In addition, if the dendrites are allowed to grow sufficiently long, they can reach the cathode, shorting the battery and causing it to catch on fire or explode.

About ten years ago Sony succeeded in utilizing graphite as an "intercalation host" on the anode side. This means that the lithium atoms are held within graphite crystals at the anode. When discharging, the lithium ions are released from the graphite and migrate toward the cathode as before, while, when recharging, the ions are inserted back into the carbon crystals and are prevented from forming dendrites. This discovery, patented by Sony, is the basis for what are now called lithium-ion cells. ViA is currently using this type of cell, manufactured by PolyStor, to power the ViA II wearable computer.

It is worth pointing out that all lithium-based batteries are bound to utilize the same anode if they are to meet the safety requirements. Thus, all such batteries are strictly "lithium-ion" by the time

they reach the market. It is conventional, however, to refer to alternatives to Sony batteries by the differentiating element, i.e. the polymer electrolyte or a new type of cathode.

Lithium-Polymer Batteries

Ultralife Batteries Inc. has developed solid-state, lithium-polymer cells that do not present safety problems as long as the charging and discharging follows the suggested guidelines. The absence of a liquid electrolyte prevents the formation of dendrites. As shown in Fig. 5.2, these cells are better than most alternatives, but not by a large margin.

With a capacity of 50000 mAh/l, Ultralife batteries offer some possibilities for implementation. For instance, the UBC3049105 battery is approximately 5 cm wide by 10 cm long (2" x 4") and only 3 mm thick. It may be possible, therefore, to bend it to a fairly tight arc and wrap it around the wrist, inside the wristband. This battery has a capacity of 800 mAh, 500 recharge cycles, and generates 3.7 V. Assuming an average power draw of 1.5 W, four of these batteries in series would generate the required 12 V and would provide 1.5 hrs of operation before recharging. If 5 cm is undesirably wide for the wristband, Ultralife is developing a new, smaller and narrower battery that will come in a variety of thicknesses and will be ready by Q4 of 1999. Two of these smaller batteries would fit on either side of the WID housing, for approximately the same capacity, and with an overall wristband width of just over 3 cm.

Li-V₂O₅ Batteries

A new high-performance lithium battery under development is based on a vanadium pentaoxide cathode. Li-V₂O₅ appears to be a strong candidate for the Wrist Interactive Device that performs better than lithium-polymer. It meets temperature requirements, has good flexibility in a thin-film format, superior energy density (with respect to both mass and volume), meets the WID's amperage discharge requirements and has very good storage characteristics.

Currently, however, Li-V₂O₅ batteries are not commercially available. The patent rights to Li-V₂O₅ rechargeable batteries are held by the Corrosion Research Center of the University of Minnesota's Chemical Engineering Department (CRC). ViA has a good relationship with this Center, including a joint effort to design mobile robot systems. This research partnership is partly funded through a DARPA STTR grant that was recently awarded to ViA, "Microrover for Tactical Land Warfare." Currently, CRC is also looking for a partner to commercialize Li-V₂O₅ technology. In the meantime, CRC does have its own limited manufacturing capabilities and is willing to fabricate custom batteries for ViA's needs.

The voltage discharge curve of these batteries is not constant. For example, a Lead-Acid cell is constant at a rate of approximately 2 V, and a Li-Ion is constant at approximately 3.5 V. As shown in Fig. 5.3, the Li-V₂O₅ curve decreases linearly from about 4 V to 0 V. Although the useable range is approximately 4 V to 2 V, the linear decrease makes the detection of the 2 V level easier. The disadvantage is that additional control circuitry is needed to ensure that the WID functions properly with this decreasing voltage. Following CRC's guidelines, ViA is

planning to develop the necessary controlling circuit to monitor the discharging and recharging process.




Fig. 5.3 Typical Voltage Discharge Curve for a CRC V_2O_5 Xero-Gel Battery (@C/4 Discharge Rate)

Among $Li-V_2O_5$ batteries, different performance characteristics can be achieved depending on how the water used in the V_2O_5 cathode manufacturing process is removed. In $Li-V_2O_5$ Xero-gel batteries the water is removed from the substrate simply by means of a vacuum pump. The high surface tension of water causes the microscopic pores of the V_2O_5 material to collapse partially while it evaporates. The result is that the overall volume of the cathode sample is approximately halved, maximizing the volumetric energy density. However, since the surface area is reduced (to roughly $100 \text{ m}^2/\text{g}$), the amperage discharge rate is relatively low.

$Li-V_2O_5$ Aero-gel formulations are at the opposite end of this spectrum. Here, the water present in the V_2O_5 cathode is first replaced with an organic solvent at atmospheric pressure by repeatedly flooding and flushing the chamber holding the V_2O_5 sample. The chamber is then sealed and brought to a high enough pressure to maintain the liquid state of the solvent while the temperature is increased to the value corresponding to the supercritical point on the P-T phase diagram. The pressure is then decreased gradually, allowing the liquid to reach the supercritical point from above. In this condition there is no distinction between the liquid and the gas phase of the solvent, causing its surface tension to vanish and allowing it to be evaporated out of the chamber and out of the V_2O_5 cathode without collapsing its pores. As a result, the reduction in volume is small and the surface area attained is quite large (approx. $450 \text{ m}^2/\text{g}$). The amperage rate is therefore maximized at the expense of volumetric energy density.

The V_2O_5 fabrication process can be modified to produce a cathode that is between these two extremes. At the present time CRC has facilities adequate for manufacturing Xero-gel V_2O_5 batteries in small quantities but not Aero-gel ones. This limitation does not hurt the development of WID prototypes since for the WID the compactness resulting from a higher energy density is a

greater premium than a higher discharge rate. CRC is nonetheless interested in working with ViA on the design of a V_2O_5 battery optimized for WID production units.

In order to reach the commercialization stage, the development of Li- V_2O_5 batteries needs to overcome several hurdles. Fig. 5.4 summarizes the main findings and some of the research planned by CRC. The discovery that motivated much of CRC's work over the last few years was that amorphous V_2O_5 holds 8 times as many Li ions than crystalline V_2O_5 , directly impacting the energy density of the resulting batteries. Because the conductivity of amorphous V_2O_5 is one thousand times smaller, CRC has adopted copper-doped amorphous V_2O_5 , bringing the conductivity up to one tenth of the crystalline oxide.

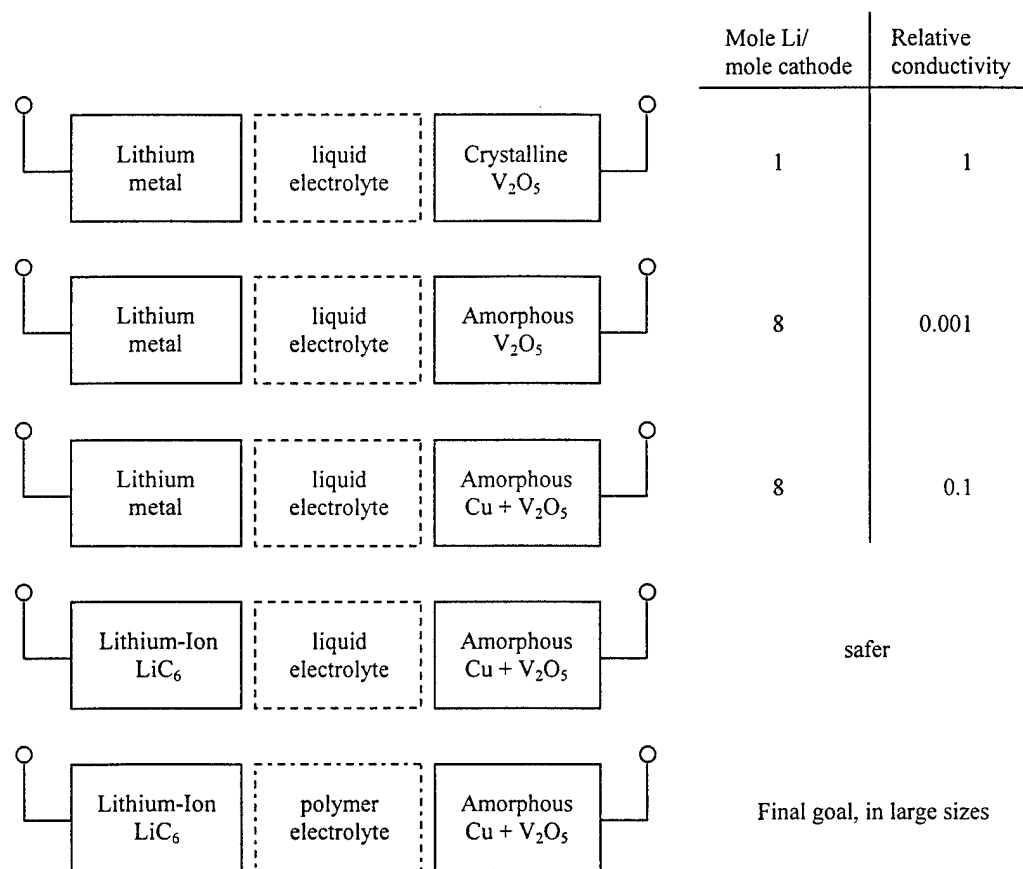


Fig. 5.4 Summary of CRC's Development of Li- V_2O_5 Batteries

It should be remarked, however, that the electron conductivity of the cathode does not have as large an impact on discharge rate as one might expect. In fact, the physical process governing discharge is, rather, ion diffusion from the electrolyte into the solid V_2O_5 cathode. In other words, ionic diffusion into the solid cathode is the "bottleneck" in the current flow of this system. Since the mean diffusion length over useful time scales depends on the physical properties of the cathode, minimizing the thickness of the solid in which the ions have to intercalate themselves optimizes the efficiency of the process—hence the advantage of

maintaining the porosity of the cathode down to microscopic scales. Thus, the ratio of electron conductivity between different versions of the V_2O_5 cathode shown in Fig. 5.4 in no way corresponds to the ratio of discharge rates. This fact makes the greater energy density of amorphous V_2O_5 a very significant advantage over the crystalline form, with only a minor decrease in discharge rate.

Currently, CRC is making batteries consistent with the second step shown in Fig. 5.4 and in sizes of roughly 200 g (7.2 oz.), which is approximately 40 times larger than a coin cell. Preliminary tests of these batteries indicate an energy density of roughly 350 Wh/Kg and 350 Wh/l. If this density can be maintained also for CRC's subsequent development steps the improvement over existing lithium batteries will be significant indeed. To put this value for energy density in perspective, we can assume an energy requirement for the WID of 4 Wh. This could be achieved by a Xero-gel battery with a volume of 11 cm³, corresponding to two prismatic cells 6 x 3 x 0.3 cm. In order to achieve the required minimum voltage (5 V) when the cells are nearly spent (2 V), a third cell would need to be added, connecting all three in series. Such cells could be mounted on the wristband of the WID. Fig. 5.5 shows the variation in specific capacity and specific energy of two of these cells as a function of cycle number.




Fig. 5.5 Variation of Specific Capacity and Energy with Cycle Number (@C/4 Discharge Rate)

Copper-doped amorphous vanadium pentaoxide has been produced in small quantities (5 g [1 oz = 28 g]) by CRC, but so far the mixing and preparation process has not scaled successfully to large quantities (50 g or more). The next task CRC has already started to undertake is to make use of lithiated carbon rather than lithium metal for the anode, thereby producing essentially a lithium-ion battery. Finally, the liquid electrolyte will be replaced by a polymer electrolyte. The final goal, therefore, is to produce a battery that combines the best features of lithium-based batteries: safe as Li-ion, of flexible design and environmentally friendly as lithium-polymer, and of greater energy density than either.

ViA is committed to taking advantage of all new technological developments that may be relevant to its products, and therefore also to the WID. For instance, a promising technology for energy storage and conversion are hydrogen fuel cells. All current fuel-cell implementations are too large to consider for the WID. However, Sandia National Laboratories is developing a microfuel cell as part of their "Fuel Cell on a Chip" program. ViA will continue to monitor the status of this effort to determine its suitability for the WID program. A review of fuel-cell technology is included in the Appendix.

Table 5.1 summarizes the main features of the battery technologies shown in Fig. 5.2.

Table 5.1 BATTERY TECHNOLOGIES

Technology	Watt-hrs/Kg	Watt-hrs/liter	Number of Recharges	Operational Temperature Range (C)	Cost	Comments	Company
Lead-Acid	30	65	500	-40 to 60	Low	3	
Nickel-Zinc	70	150	600		Medium		
Nickel-Cadmium	40	115	200	-40 to 70	Medium	1, 2, 3	
Nickel-Metal Hydride	50	170	300	-20 to 45	High	2, 3	
Zinc-Silver	200	400	Low		High		
Zinc-Air	150	140					
Zinc-Manganese	75	200	25		Low		
Lithium-Ion	110	250	1200	-20 to 60	High	5	MoliCel, Polystor, Sony
Lithium-Polymer	130	210	500	-20 to 60	High		Ultralife
Lithium-V ₂ O ₅ : Aerogel	560	appx. 350	500	10 to 60		5, 6	
Lithium-V ₂ O ₅ : Xerogel	420	appx. 400	500	10 to 60		5, 6	
Fuel-Cell			>10000		High	4	ONSI Corp., Ballard Power Systems

1. Number of recharges is strongly influenced by the discharge/recharge cycle.

Ni-Cads will form a "memory" of this cycle so that, if they are not fully discharged, the battery life is shortened.

2. Poor shelf life characteristic. Appx. 1-2% of remaining energy is dissipated each day at room temp (the higher the temp, the greater the dissipation).

3. Disposal problems -- hazardous material.

4. Securing a safe supply of hydrogen for a portable applications is an issue -- methanol and hydrides (bulky and heavy) are two approaches.

5. Thin film, flexible form factor.

6. These numbers apply only to coin-size lab samples. Performance of production units of greater sizes should be somewhat lower than these figures.


Appendix

Review Of Fuel Cell Technology

Fuel cells are an efficient, nonpolluting power source that produce no noise and have no moving parts. They have been providing electricity on spacecraft since the 1960s, more recently for electricity-generating plants, and will soon be used as a power source for automobiles.

Fuel cells, unlike batteries, are almost endlessly rechargeable. The cells run on hydrogen, which reacts with oxygen from the air to generate a voltage between two electrodes; the reactions occur in a chemical mediator known as an electrolyte. Compared with conventional fossil-fuel power sources, fuel cells are exceptionally clean and efficient. Their primary waste product is water; natural gas-fueled cells do produce some carbon dioxide as well, though less than would be created if the fuel were burned.

Currently, ONSI Corporation in Windsor, Conn., a subsidiary of International Fuel Cells, is the largest U.S. commercial manufacturer of fuel cells. Seventy-four of its units, each the size of a minivan, are now in operation in locations such as hospitals and remote hotels where grid power is expensive and reliability is worth a premium. (An ONSI installation in Groton, Conn., is consuming methane from a landfill, thereby both generating power and siphoning off an explosive waste gas; the U.S. Department of Energy is supporting a similar project.) Each ONSI cell provides 200 kilowatts of power. ONSI's marketing manager, Gregory J. Sandelli, states that in 1.25 million hours of total use, his company's cells have remained in operation 95 percent of the time--a figure that bests on-site, diesel-powered generators. The units, which use phosphoric acid as an electrolyte, are designed to last 20 years.



Phosphoric acid-based cells tend to be heavy, which makes them less than ideal for use in vehicles. Other companies are developing cells that are specifically designed for that purpose. Ballard Power Systems in Vancouver, B.C., has been developing a lighter fuel cell, the "proton-exchange membrane" (PEM) type. In place of phosphoric acid, PEM cells employ a thin polymer membrane as their electrolyte. Ballard is planning to launch pilot fleets of a fuel cell-powered passenger bus in Vancouver and Chicago. Both Toyota and Daimler-Benz (the latter has been collaborating with Ballard) have exhibited PEM cell-powered automobiles. The main difficulty with using fuel cells in passenger vehicles is in supplying the hydrogen. Although the gas can be compressed, companies are worried about safety should a pressurized hydrogen tank be damaged in an accident (e.g., remember the Hindenberg). Toyota's demonstration car solved the problem by storing the hydrogen in hydrides, compounds that can stockpile hydrogen in a safe, but heavy and cumbersome, form. Toyota and Daimler-Benz are also exploring a technique for producing the hydrogen as needed by breaking down a fuel such as methanol, which is far easier to transport and pump into a car. In August, 1998, Royal Dutch Shell demonstrated their

"Millenium Taxi." The taxi uses a 5 kW alkaline fuel cell, built by Zevco, that charges the car's batteries which in turn provide power to the electric motor. Hydrogen is stored at 3000 psi. The current cost of ZEVCO's fuel cell is \$2,000, but the company hopes to bring that figure down to around \$300 by 2001 and to \$60 in ten years.

When portability is not an issue, as it is not for municipal utilities, other possibilities open up. "Molten carbonate" and "solid oxide" fuel-cell technologies, for instance, could bring extraordinary efficiencies to power-generating stations. These devices run at far higher temperatures than PEM or phosphoric acid cells. Molten carbonate and solid oxide cells might be able to achieve an impressive efficiency of 55 percent or more, according to Robert R. Rose of Fuel Cells 2000, an advocacy group. The hot steam and carbon dioxide they produce can be used to drive a gas turbine that generates additional electricity, an approach that could push their efficiency to an unheard-of 80 percent. They are potentially up to twice as efficient as a typical oil- or coal-fired plant.

Energy Research Corporation (ERC) in Danbury, Conn., has built a two-megawatt pilot molten carbonate plant for the municipality of Santa Clara, Calif. The fuel cells run at about 650 degrees Celsius (1,200 Fahrenheit) and consume hydrocarbon fuel that is re-formed into hydrogen right inside the cells. Problems with electrical shorts within the cells have restricted the plant's output so far to about one megawatt. As the ERC's William V. Baker is at pains to point out, however, that still makes the Santa Clara plant the highest-powered fuel cell unit in the U.S. The company is currently developing a modified design rated at 2.85 megawatts.

Critics of the molten carbonate technology note that such cells use as electrolytes highly corrosive molten salts that complicate design and maintenance. Westinghouse is therefore focusing on the rival solid-oxide approach. These cells, which employ solid metal oxides as their electrolyte, operate at even higher temperatures than molten carbonate, close to 1,000 degrees C (1,800 F). Nevertheless, Christian A. Forbes, manager of business development for solid-oxide fuel cells at Westinghouse, states that some component cells have logged 69,000 hours (almost eight years) of activity. Westinghouse's cells have an open-ended tubular design that makes them look something like a church organ. That approach minimizes sealing problems caused when materials expand as the temperature rises. The company is now installing a 100-kilowatt test unit in the Netherlands. A second unit may follow within a year, and a decision on whether to initiate full-scale production is expected in 1999.

Ceramatec of Utah is one of several companies developing solid-oxide fuel cells in an alternative "planar" configuration that incorporates stacks of disk-shaped electrodes. Planar designs traditionally suffer from sealing problems, but Steven Visco and his colleagues at Lawrence Berkeley National Laboratory think they have identified a solution. Visco has found a way to make thinner layers of the electrolyte, which allows the cells to run at a slightly cooler 800 degrees C (1,500 F). That change will eliminate the sealing problems and double the power output, the company hopes. Ceramatec is building a high-efficiency cogeneration unit in collaboration with Sulzer in Switzerland.

Chapter VI

DISPLAY AND COMMUNICATION DRIVERS

ViA has started the development of the drivers to interface to the communication, display, and audio hardware. The emphasis so far has been on the display drivers and the beginnings of the communication drivers. Fig. 6.1 shows a block diagram of the distributed video system. The main components of the software in this system are now discussed.

Drivers on the Host Computer (ViA II)

The Windows architecture requires applications to request system services at a high level through API (Application Programming Interface) calls, which requires applications to use well-defined software function calls. If an application wants to draw a circle, it only has to tell windows to draw a circle because the capability is defined as part of the interface. There are two basic subsystems that Windows uses to handle the interface. The GDI (Graphics Display Interface) is responsible for routing calls from a Windows application. There is also a DIB (Device-Independent Bitmap) subsystem that is responsible for generating many of the actual images requested by GDI. When a display driver loads, it communicates with GDI to tell GDI what capabilities it has. GDI uses its knowledge of those abilities to determine how to handle application requests. The request for a circle can be routed by GDI to the display driver if the display hardware can do the job, or GDI can pass the task to the DIB engine. If the request goes to the DIB engine, DIB draws the circle in its own way. Perhaps the circle is drawn as a series of short, straight lines; perhaps DIB calculates the image pixel by pixel. The point is that the operating system can decide whether to handle much of the work itself, or pass the request on to a display driver.

The bare-bones VGA driver actually has a very limited capability. It can basically draw a scan line (a raster line) or a line or series of lines between arbitrary points. It can draw a string of text. It can also draw bitmaps (a screen icon is an example of a bitmap) of various flavors, or write to a single pixel on the screen. This is not to trivialize the work the driver does. It has to handle many operations that combine or modify images, such as highlighting text, and many of these operations are convoluted as the driver deals with memory that may be in different segments or cross byte boundaries. It has to handle a mouse cursor. A modern driver will use hardware to handle more complex calls, such as the example of the circle.

In the present implementation of the distributed video system, the VGA driver has been extended by adding routines to gather data and then write the results over a serial port, to the WID. In some cases, such as the line drawing routine, the data are transferred at the entry point in the VGA driver. In other words, some of the call parameters are passed as they arrive at the driver. If the parameter is merely the address of data, the data themselves are gathered and transmitted.

Strings are handled slightly differently. The string function is called a number of times without displaying anything on the screen. Data are gathered and sent to the WID at the point that the driver finally displays the string.

Windows-Based Host PC
(Via II)

WID

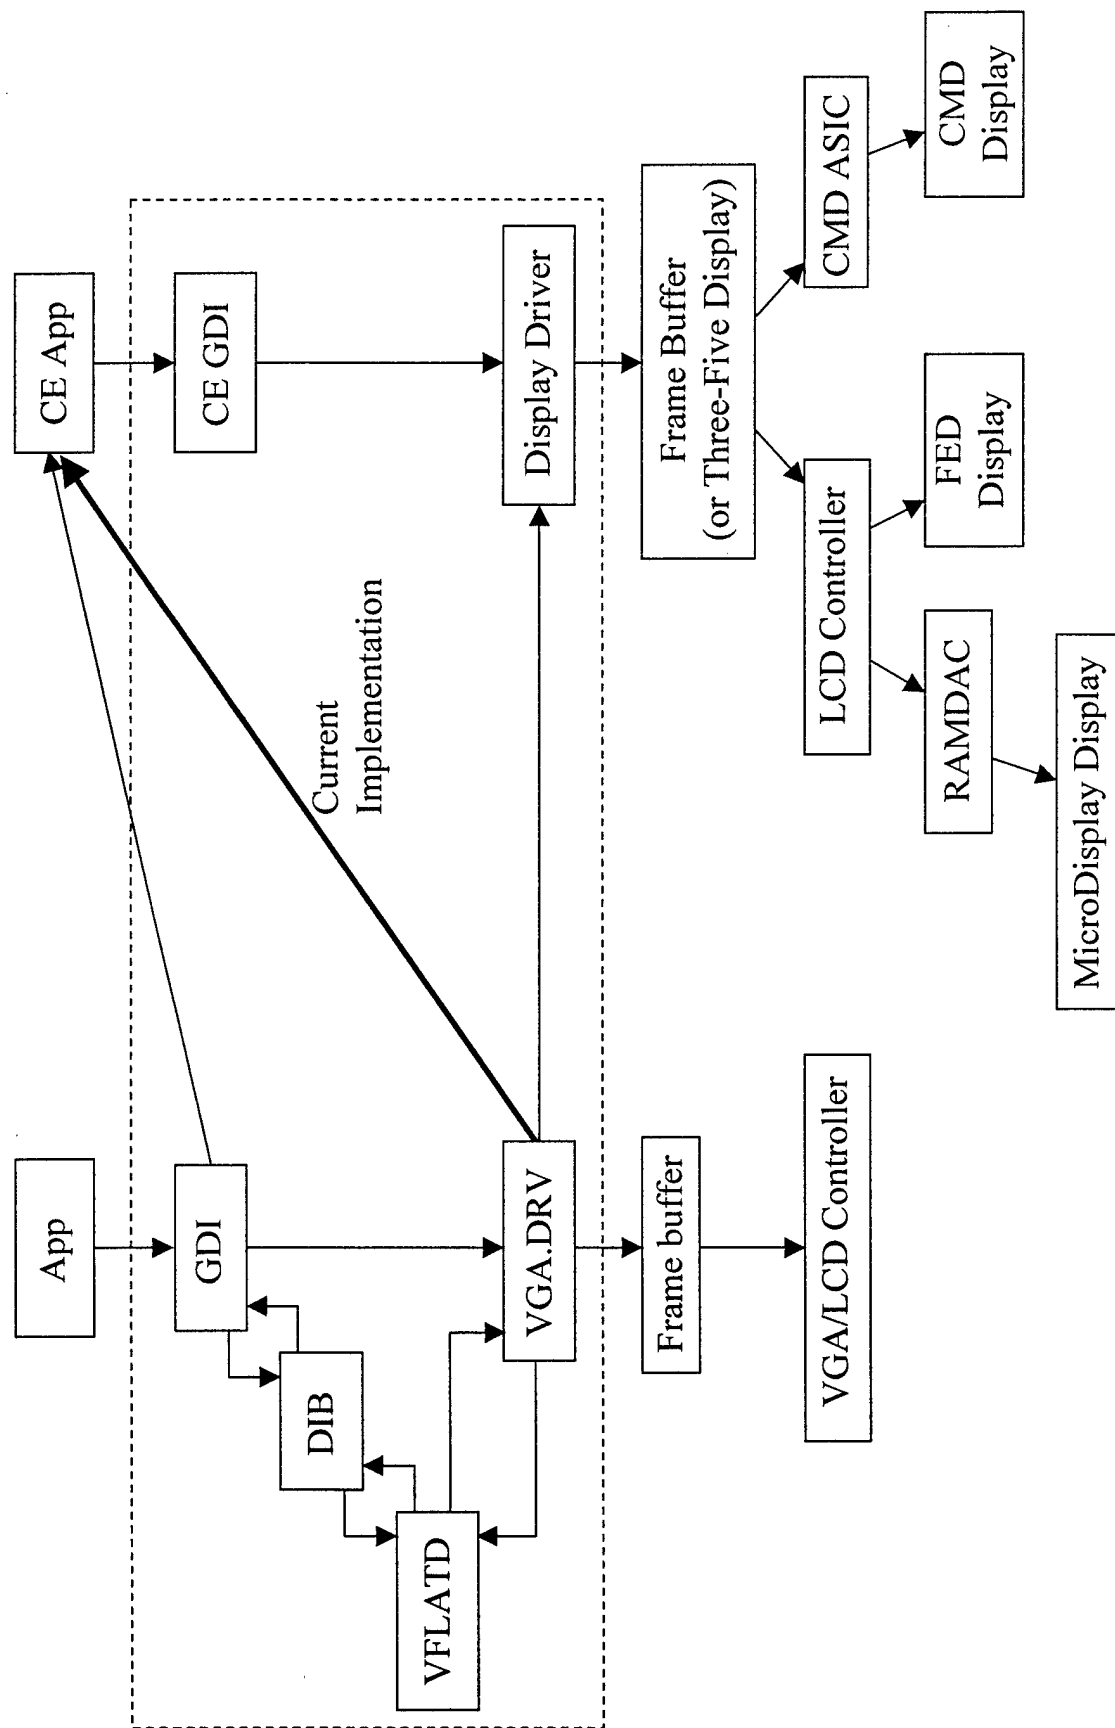


Fig. 6.1 Distributed Video System for the WID

Bitmaps are handled with a third variation. The routine examines the parameters to determine how to handle any of a number of different cases. At this time, only two branches of the decision tree are implemented. One case is when the driver realizes that the bitmap is a rectangle of a solid color. In this case, the coordinates and the color of the rectangle are passed across the serial line. The other implemented case is the transfer of a bitmap from system memory to the display. In this case, the bitmap is reconstructed from the system memory and sent pixel by pixel.

These implemented functions are surprisingly powerful. Of particular note is that a window frame is drawn entirely using these functions. Windows creates the buttons as system memory icons, as are most other icons. All of the rest of the frame is made of solid color rectangles. Even the lines represent degenerate cases of a rectangle with either no height or width.

High-Level CE Application on the WID

A Windows CE application has been developed to run on the SideARM board, the WID board, or any other platform running CE. ViA is utilizing a palmtop made by Hewlett-Packard, the Jornada, as an initial test platform for high-level CE software. The Jornada is built around the SA1100 microprocessor and the SA1101 support chip, so it serves as a very good test platform for the WID.

The CE application ViA developed retrieves data coming over the serial port, interprets the packets as containing data or graphics commands, and calls CE GDI routines to render the graphical information. A two-character word precedes each structure being sent and identifies what type of graphic should be displayed (bm for bitmap, pl for polyline, etc.). The identifier is followed by the size of the structure in bytes, which is used to read the remaining data. Each type of graphic is handled differently. For example, a polyline is an array of points. The points are simply drawn on the screen, connected by the current pen. A bitmap is more complex. It is made up of an X and Y origin, an X and Y size, and an array of palette references. The palette references are converted to RGB values and the corresponding pixel is set to that color. Each graphic's structure is defined in the video driver documentation (Datastruc.doc).

Porting Windows CE to the SA1100 Microprocessor

Porting Windows CE to a new hardware platform is a multiple stage process. Microsoft's Windows CE Embedded Toolkit version 2.10 separates the development process into 2 directories, *platform* and *public*. The platform directory contains all the hardware-specific code. It contains the code to initialize the specific CPU, and how to talk to the peripherals on the development platform. The public directory specifies what components of Windows CE to include in the binary image, i.e. whether or not to use mouse & keyboard drivers, whether or not to include serial port support, etc. It is also the place to include any custom programs into the binary image.

The starting point has been to port CE to the development board for the SA1100, the WID's microprocessor. Although such a port had already been done by Intel, it was for their previous development board, the Brutus board. The current development board, the SideARM, was

released recently and Intel has not yet completed the CE port to it. Since the Brutus board had been optimized for flexibility at the expense of performance, ViA elected to utilize the SideARM and its companion board, the SideKick, in spite of the greater effort that the OS development required. As discussed more fully in Chapter 3, the WID Development Board #1 replicates the architecture of the SideARM and SideKick boards.

Microsoft supplies sample platform and public directories, so the first step in porting Windows CE is to create a copy of the platform directory that is most similar to the target hardware platform. Intel has released source code for the Brutus development board, so a copy of the Brutus platform directory was used to begin a port to the SideARM. Several components in this directory need to be modified for Windows CE to run on the SideARM board:

1. **Boot loader.** The boot loader is a piece of software running on the target platform which initializes the CPU and transfers the Windows CE binary image from the development workstation to the target platform. The SideARM board came with the Angel Debugger. Using this debug interface, it is possible to transfer a binary image across a serial cable. The Angel Debugger is therefore being used as the boot loader at the current time.
2. **Kernel Initialization.** After the Windows CE binary image has been moved to the target, the boot loader then transfers execution to the Windows CE kernel. At kernel start-up, a number of things must be initialized: ports, timers, real-time clock, interrupt requests, mapping of ISRs to hardware IRQs, checking for extension DRAM, and implementation of power management functions.
3. **Drivers.** Once the kernel initialization code is functioning properly, device drivers need to be developed to use the peripheral hardware. Specifically, display drivers and serial communication drivers need to be implemented for the WID project.

After the hardware-specific code has been successfully developed, a specialized implementation of Windows CE needs to be chosen. Microsoft has split the operating system into chunks called *modules*. Modules vary in function from hard disk support, serial communication support, PCMCIA card support, audio components, to GUI components. When choosing a Windows CE configuration, each module can be included or discarded depending on what is desired in the operating system. Obviously, the more modules included in the OS, the larger the binary image will be. The WID's remote display application specifically requires the graphical interface and serial communication support. These modules have been selected and, to conserve RAM and ROM space, most of the other OS modules have been excluded in our configuration. The final step in porting Windows CE is to write and include ViA's custom software. The remote display application needs to be written, compiled into the binary image of the OS, and then automatically executed when the operating system boots.

Once CE has been ported to the SideARM board, it will be straightforward to port it to the WID board. In fact, the WID Development Board #1 was specifically designed to reflect the configuration and layout of the SideARM to facilitate the porting of the software. As future WID development boards are shrunk to increasingly smaller sizes, the OS and drivers will be optimized incrementally to keep up with the changes in the hardware. However, it is expected that these changes will have a minimal effect on the OS and drivers architecture.

Fig. 6.2 shows an approximate roadmap of the development of the drivers for the WID.

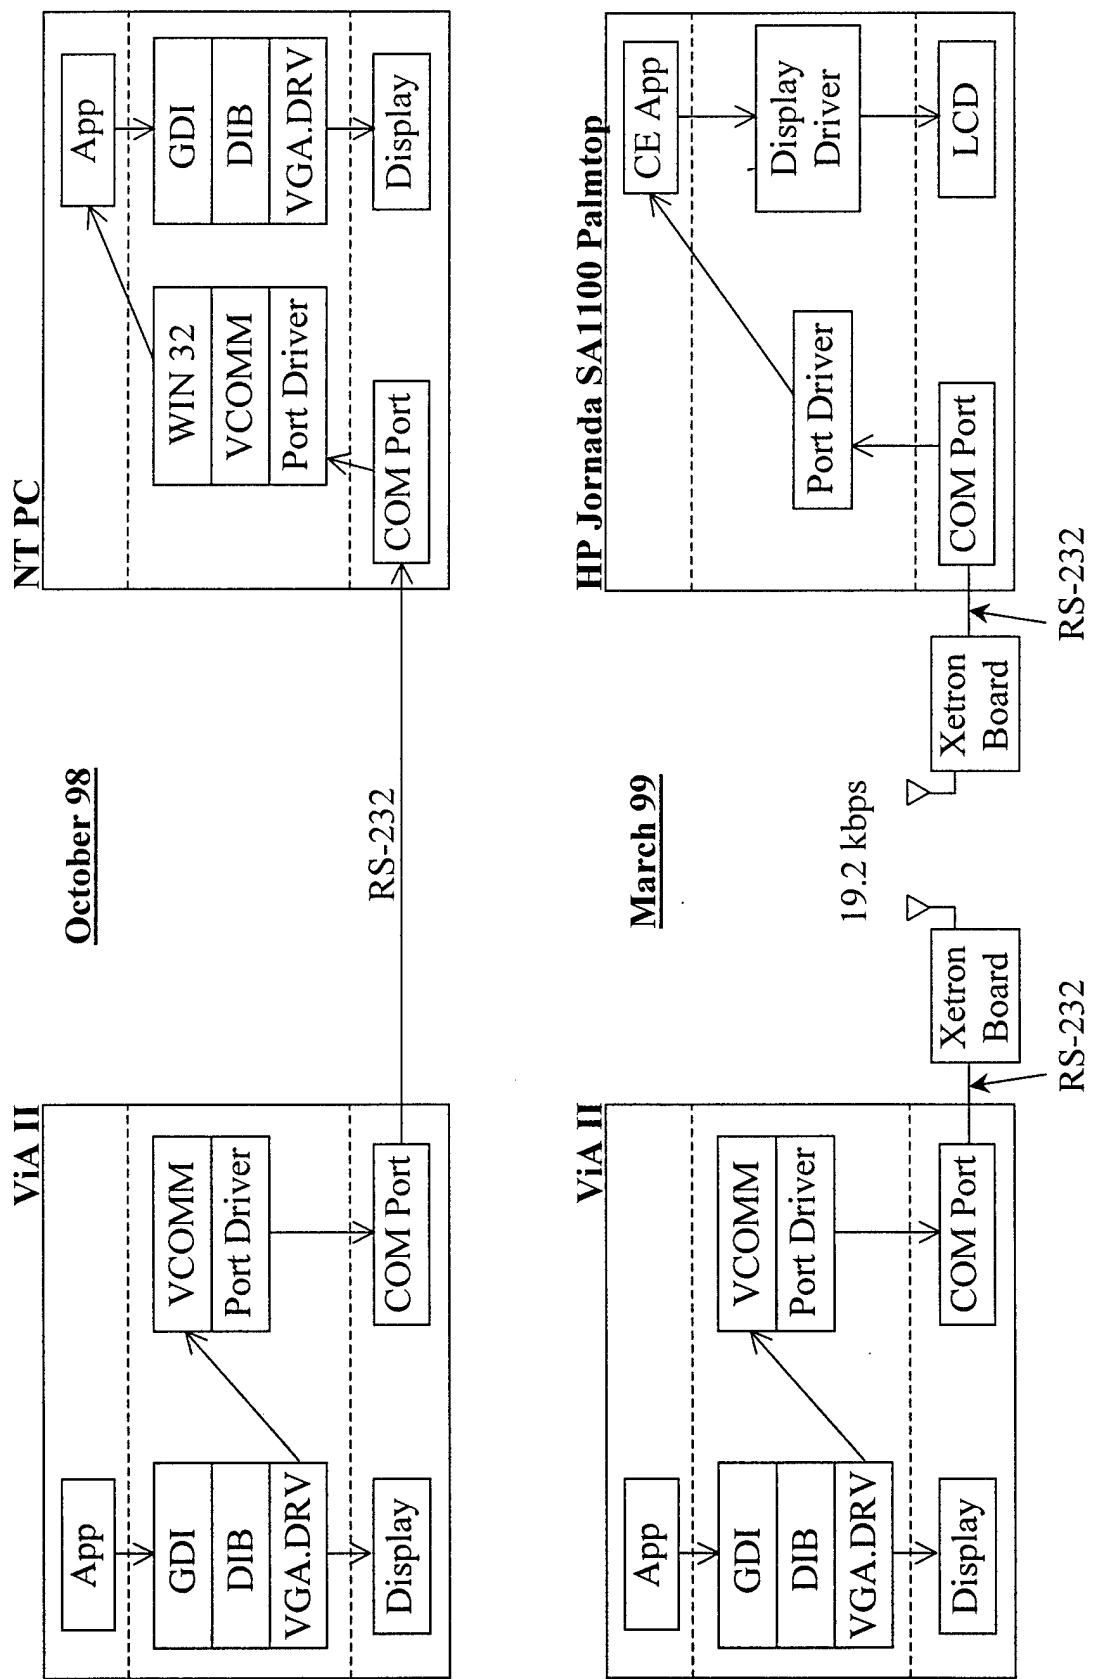


Fig. 6.2 WID Driver Development Roadmap




Fig. 6.2 WID Driver Development Roadmap (Continued)

Chapter VII

DEVELOPMENT OF WID EVALUATION BOARD #2

Overview

Fig. 7.1 shows the planned WID Development Board #2. As for the first board, the main purpose of this board is to facilitate the development of the software required for the WID and to test hardware components. The emphasis in this board will be the audio system, as discussed more fully later in this chapter. Concurrently, the CMD display will be integrated with the optics designed by DisplayWear in the form of a more sophisticated physical model. This new model will have a parallel digital interface to a connector on the board to allow the SA1100 to drive the display as a memory-mapped device. The data and power signals will be carried by a cable that will plug into the side of the WID physical/optics model.

In addition to the CMD display, this board will be able to drive the FED display, the MicroDisplay display (if it becomes available), and the Three-Five display through the same memory-mapped interface utilized by the CMD display. At the time of this writing it appears that the CMD display is the most likely to be ready for integration with the WID.

Xetron has stated it will deliver to ViA the ASIC version of their Hornet radio boards in time for incorporation into the WID Development Board #2. In this new version of their radio the digital and RF sections are reduced to one chip each. The digital ASIC will present a parallel interface to the SA1100, which will perform a number of tasks now serviced by an Atmel 4 MHz microcontroller. As this Atmel microcontroller is a bottle-neck in the current design, the new radio board will allow the bit rate to increase to 115 kbps. However, a significant effort will have to be expended to port the software currently running on the Atmel processor to the SA1100 and the system may not become demonstrable until Q3 99.

ViA will keep monitoring the availability of the Bluetooth radio. If chip-scale packaging for a Bluetooth radio becomes available before October 99 it will be incorporated into the WID design, otherwise ViA will keep relying on the Xetron radio. Although slower in bit rate, the two-chip Xetron solution is within range of the correct form factor to fit inside the WID.

The results of the testing of WID Board #2 will be used to design WID board #3. This will be the first board for which a specific effort will be made to shrink its physical size. In fact, by this point the optimum hardware components for each WID subsystem will have been selected. The design of this third board will be done in close collaboration with a mechanical engineer in order to address the relevant packaging issues. Concurrently with these developments, the batteries for the WID (probably Ultralife) will be tested and the control/power supply circuit will be designed in order to allow Board #3 to run off the batteries.

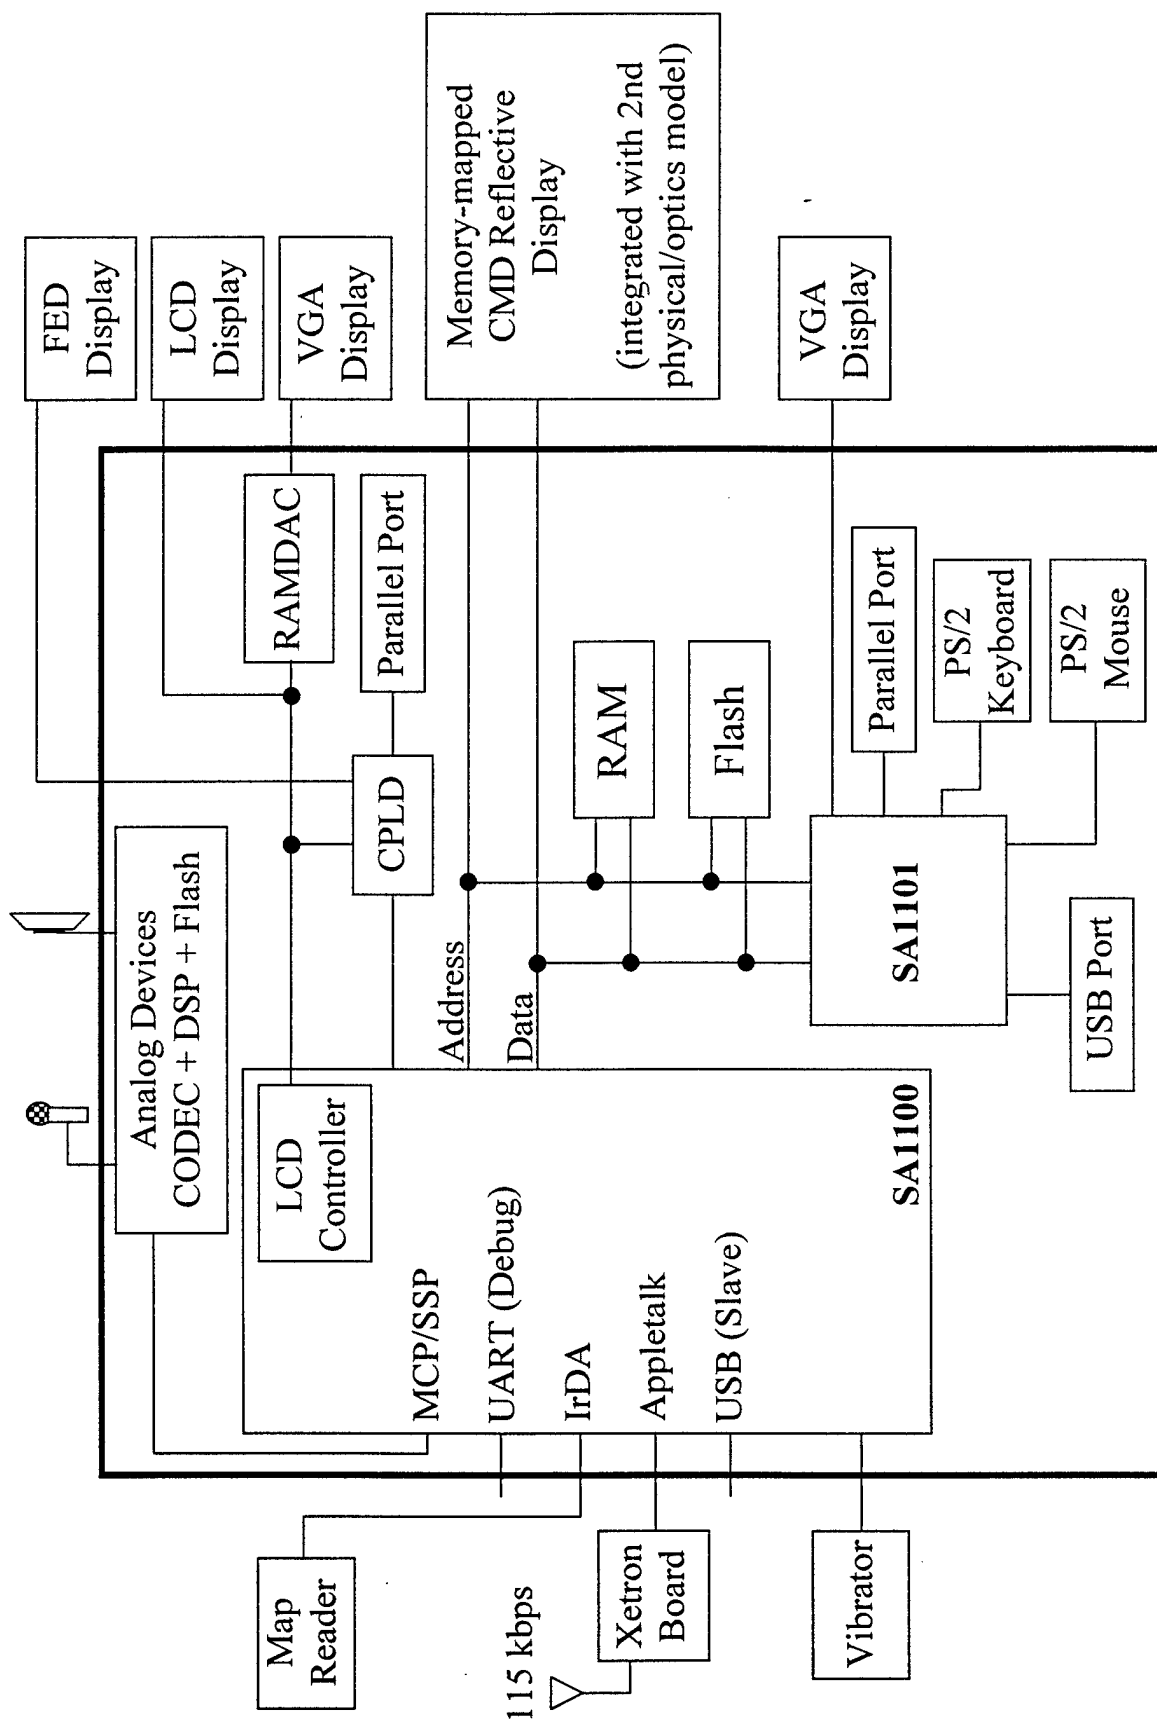


Fig. 7.1 Planned WID Development Board #2

Audio System Characteristics and Implementation Considerations

The WID's CODEC is the heart of the audio interface with the user. Since voice commands constitute the pointing device for the WID platform, accurate digitization is important to maximize the recognition rate of the voice-recognition engine. In the other direction, the computer's response must be accurately rendered on the WID's speaker. This section discusses the main issues related to these tasks; namely,

- CODEC, microphone, and speaker selection
- Filtering
- System hook-up
- Interface to the microprocessor
- Voice compression and noise cancellation

CODEC, Microphone, and Speaker Selection

The required performance characteristics of the CODEC are: 11 kHz sampling rate, 16 bits of vertical resolution, programmable input gain, and 3 V for low power consumption. ViA has identified two chips that satisfy these requirements, as shown in the following table:

Table 7.1 Comparison of Two CODECs			
CODEC		Analog Devices AD73311L	Philips UCB1200
Sampling Rate	(kHz)	(programmable)	(programmable)
Theoretical Vertical Resolution	(bits)	16	12
Actual Vertical Resolution	(bits)	12	8
Theoretical Dynamic Range	(dB)	96	72
Actual Dynamic Range	(dB)	70	50
Programmable Input Gain	(dB)	38	22.5
Power Supply Voltage	(V)	3	3
Power Consumption	(mW)	60	80

In both cases the actual dynamic range is smaller than the theoretical. The theoretical range could have been attained by designing more extensive digital filtering into the chips, but that would have led to greater cost, size, and power requirements for these parts. In the present case, the differences in dynamic range are not an issue since the critical parameter is the sensitivity of the microphone. Both microphones ViA has examined have a sensitivity of approximately 50 dB, thereby rendering any additional resolution on the part of the CODEC useless. In this respect, therefore, the two CODECs are roughly equivalent.

The performance characteristics of two microphones ViA has examined are as follows:

Table 7.2 Comparison of Two Microphones

Microphone		Andrea CMS-27	NCT (Siemens) SMM
Sensitivity	(dB)	50	54
Power Supply Voltage	(V)	2	1.5-6
Directionality		Omni	Omni
Output Impedance	(k Ω)	2.2	2
Current draw	(μ A)	500	50
Frequency Response	(Hz)	30-4,000	20-10,000
Noise canceling		Yes	No
Size	(mm)	6 dia. x 2.7	3 x 3

The two CODECs examined are similar also in how they handle the filtering of high-frequency noise. Both employ a sigma-delta modulator as the heart of their A/D converter. This type of converter benefits from a simple design and relies on a sampling rate many times higher than the highest resolution desired to “push” any aliased high-frequency noise away from the band of interest. Subsequent decimators and low-pass FIR filters produce a clean digital representation of the voice-band microphone signal that is passed on to the microprocessor for further processing.

Whereas the AD73311L is clearly a better CODEC, the StrongARM’s serial port to the CODEC is designed and tested to support the UCB1200’s interface and protocol. In addition, the evaluation board for the SA1100, the SideARM, comes with the Philips CODEC. ViA will therefore start the development using the UCB1200 CODEC and later in the project it will switch to the AD73311L. The difficulty of interfacing the latter to the SA1100 is alleviated by a new chip being developed by Analog Devices. In fact, the same SA1100 serial port that supports communication with the UCB1200 can be set to a different mode to act as a synchronized serial port (SSP). In this mode it can communicate with the Analog Devices ADSP2185L DSP.

Analog Devices has completed the development of a new chip, the AD73522, that combines three different dies in the same PBGA package: the AD73322L CODEC, the ADSP2185L DSP, and 64 kB of Flash memory. The AD73322L implements two AD73311L CODECs in the same die. Because the audio data may have to be processed for noise and echo canceling and will then have to be compressed before it is sent over the radio link, some of these functions could be performed by the DSP in order to relieve the SA1100 processor. This new chip would place a DSP that supports the SSP protocol between the SA1100 and the AD CODEC, thus eliminating the interfacing obstacles.

The final configuration will depend also on the available voice processing software and on which products will be easiest to port to the different chips. At the present time ViA is examining software developed by L&H, Voxware, and Noise Cancellation Technologies.

Fig. 7.2 shows a block diagram of possible configurations.

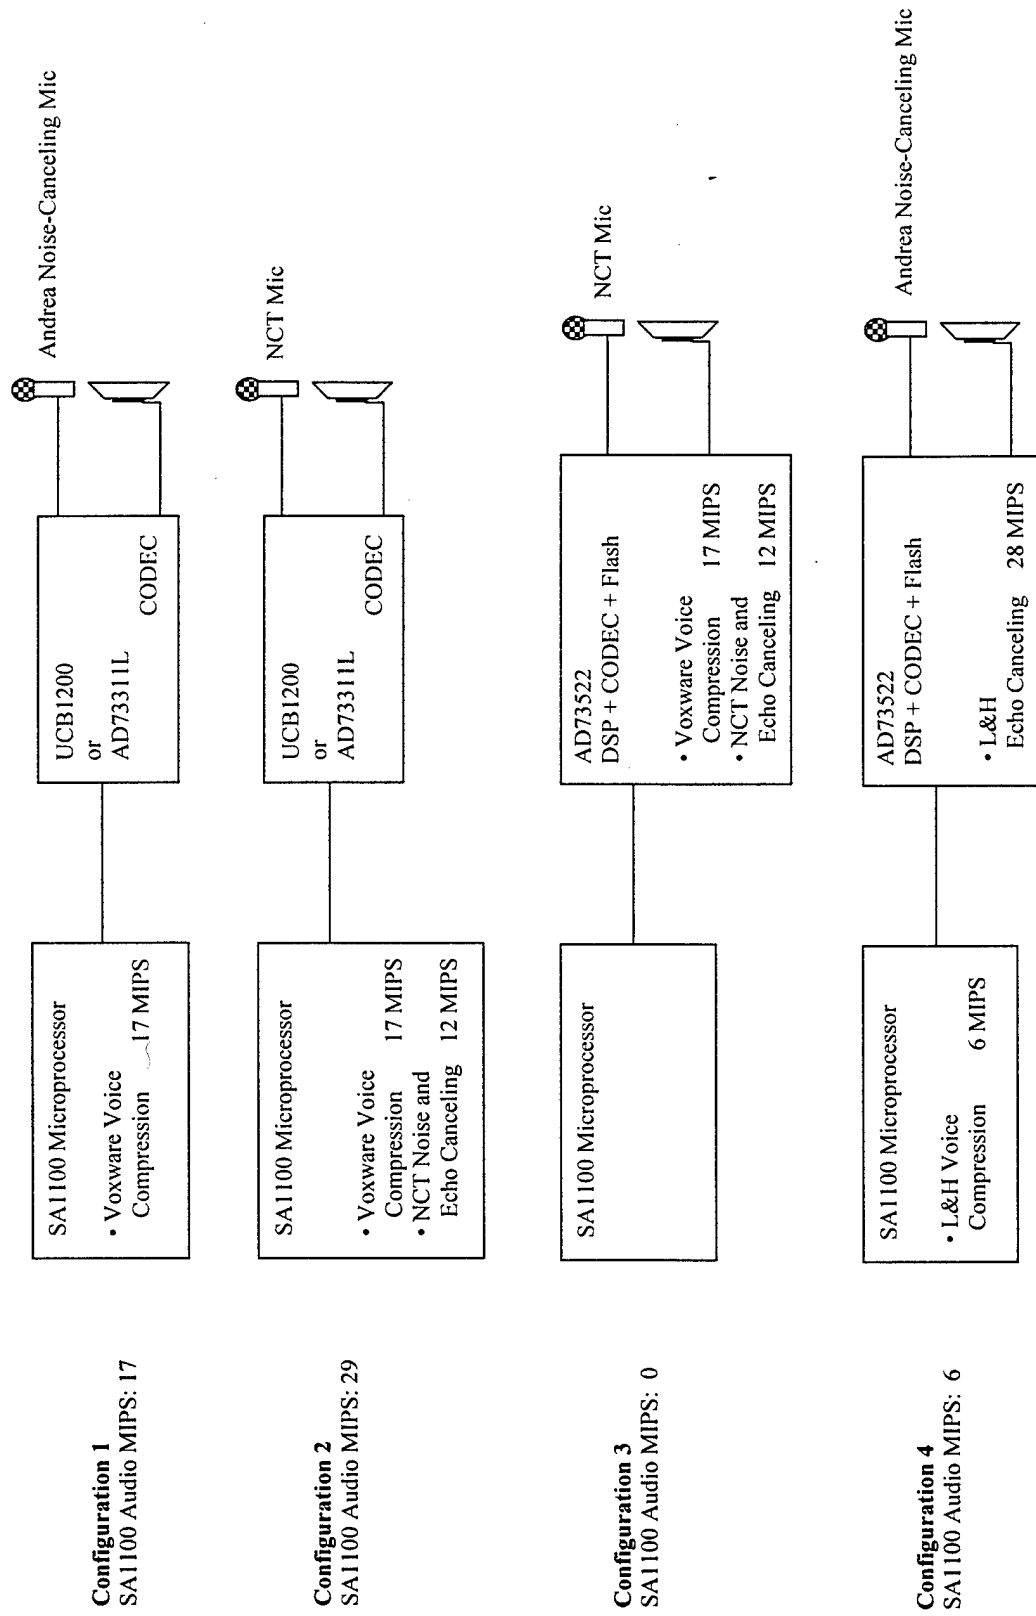


Fig. 7.2 Possible Configurations of Audio System

Filtering

The over-sampling and digital filtering performed by both CODECs are sufficient to handle the majority of noise problems that might be present in the audio signal. However, the data sheet for the Analog Devices CODEC advises to use a single-pole, low-pass, passive RC filter before the CODEC as an additional safeguard. This section compares the frequency response of such a passive filter to that of an active filter and a filter/pre-amplifier.

Low-Pass RC Filter

The following is a simple RC filter:




Fig. 7.3 Single-Pole, Low-Pass RC Filter

The in-phase and quadrature components of the response to a sinusoidal input are given by

$$\begin{aligned}\frac{A}{V_i}(f) &= \frac{1}{1 + 4\pi^2 f^2 R^2 C^2} \\ \frac{B}{V_i}(f) &= \frac{2\pi f R C}{1 + 4\pi^2 f^2 R^2 C^2}\end{aligned}\tag{3.1}$$

where f is in Hz, V_i is the amplitude of the input signal $V_i \cos \omega t$, and ω is the radian frequency.

Low-Pass Active Filter

The filter shown in Fig. 7.4 can be tuned to a desired cutoff frequency and roll-off rate. The in-phase and quadrature amplitude components of the response are given by, respectively,

$$\begin{aligned}\frac{A}{V_i}(f) &= 4\pi^2 f_0^2 \left[\frac{f_0^2 - f^2}{\Gamma^2 f^2 + 4\pi^2 (f_0^2 - f^2)^2} \right] \\ \frac{B}{V_i}(f) &= 2\pi f_0^2 \left[\frac{\Gamma f}{\Gamma^2 f^2 + 4\pi^2 (f_0^2 - f^2)^2} \right]\end{aligned}\tag{3.2}$$

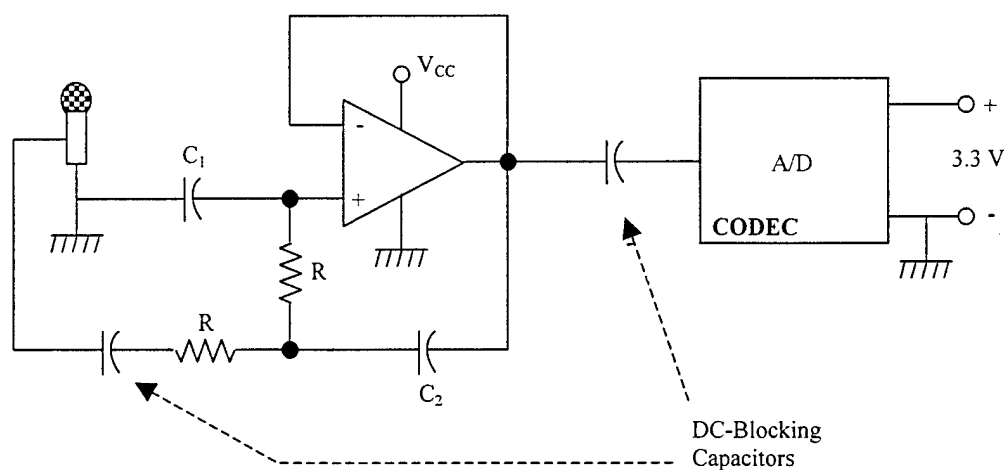


Fig. 7.4 Low-Pass Active Filter

where Γ is the damping constant:

$$f = \frac{\omega}{2\pi}, \quad f_0 = \frac{1}{2\pi R \sqrt{C_1 C_2}} \quad \text{and} \quad \Gamma = \frac{2}{RC_2} \quad (3.3)$$

Fig. 7.5 shows a comparison of the amplitude and phase response curves of this filter and the single-pole RC filter. In spite of the slower roll-off, the RC filter is likely to be sufficient.




Fig. 7.5 Comparison of Frequency Response of Active and RC Filters
 (Active Filter: Solid line, $R = 5 \text{ k}\Omega$, $C_1 = 1 \text{ nF}$, $C_2 = 2.5 \text{ nF}$, $f_0 = 20132 \text{ Hz}$
 RC Filter: Dot-dash line, $R = 10 \text{ k}\Omega$, $C = 1 \text{ nF}$)

Low-Pass Pre-Amplifier Filter

The filter shown in Fig. 7.6 amplifies the signal at low frequencies and approaches a gain of 1 as $f \rightarrow \infty$. The amplitude response is

$$\begin{aligned} \frac{A}{V_i}(f) &= \frac{4\pi^2 R_2^2 C^2 f^2 + (1 + R_2 / R_1)}{4\pi^2 R_2^2 C^2 f^2 + 1} \\ \frac{B}{V_i}(f) &= \frac{2\pi R_2 C f (R_2 / R_1)}{4\pi^2 R_2^2 C^2 f^2 + 1} \end{aligned} \quad (3.4)$$

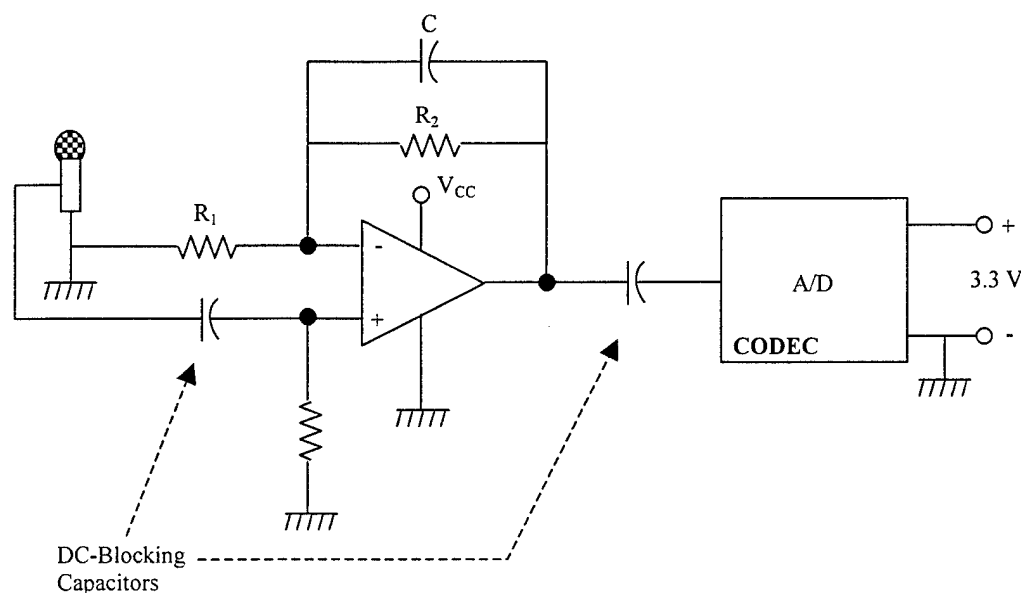


Fig. 7.6 Low-Pass, Pre-Amplifier Filter

The gain at DC conditions is given by $G = 1 + R_2 / R_1$. The amplification is performed partly by this circuit and partly by the CODEC. Fig. 7.7 shows typical frequency response curves.




Fig. 7.7 Frequency Response of Low-Pass Filter/Pre-Amplifier
($R_1 = 10 \Omega$, $R_2 = 1 \text{ k}\Omega$, and $C = 10 \text{ nF}$)

System Hook-Up

Both the Andrea and the NCT microphones are electret, or condenser, microphones and need a DC power supply. The sound pressure is translated into a voltage signal connected to the base of a transistor, which acts as the first stage of the amplification process by tapping the DC power supply. Fig. 7.8 shows the hook-up of the audio system if the Philips CODEC is utilized.

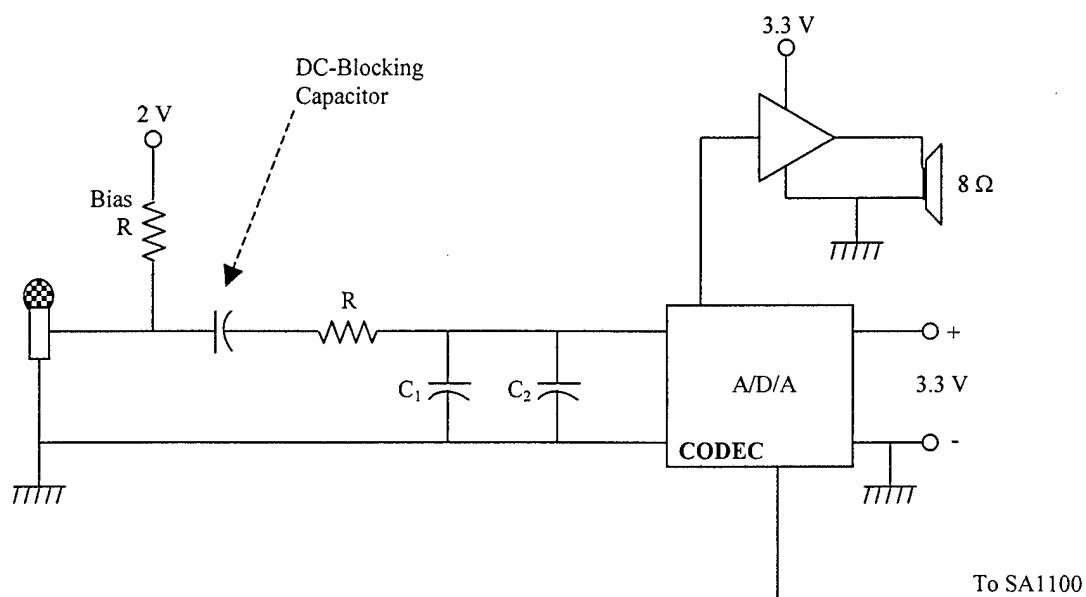


Fig. 7.8 Audio System Hook-Up with Philips UCB1200 CODEC

Two capacitors are needed to achieve adequate attenuation of high-frequency noise over at least 5 decades: when C_1 starts breaking down and behaving like an inductor, C_2 takes over and absorbs the higher-frequency noise. In order to drive the $8\ \Omega$ speaker an amplifier such as the Philips TDA1519A is needed.

The AD73322 CODEC integrated into the AD73522 chip includes an op-amp stage at the analog input that is not present in the AD73311L. This op-amp can be used to provide an active filter very similar to that of Fig. 7.4, the difference being that here it is inverting. Fig. 7.9 shows the microphone hook-up. In this figure, C_2 is the DC-blocking capacitor and the external op-amp provides clean, stable power at an easily set voltage. The frequency response of this filter/pre-amplifier is qualitatively very similar to that shown in Fig. 7.5.




Fig. 7.9 Microphone Hook-Up with AD75322 Integrated Audio Chip

Interface to the Microprocessor

The AD75322 interfaces to a host processor via a synchronous serial port on the DSP. One of the serial ports on the SA1100 can be configured either as a Multimedia Communication Port (MCP) or as a Synchronized Serial Port (SSP). The latter allows several settings that make it compatible with the DSP's serial port (SPORT0).

During development the same serial port can be used to download the program to be executed by the DSP into the flash memory. Ultimately, ViA will implement this procedure using the radio link, so as to avoid taking apart the hardware when new versions of the software become available.

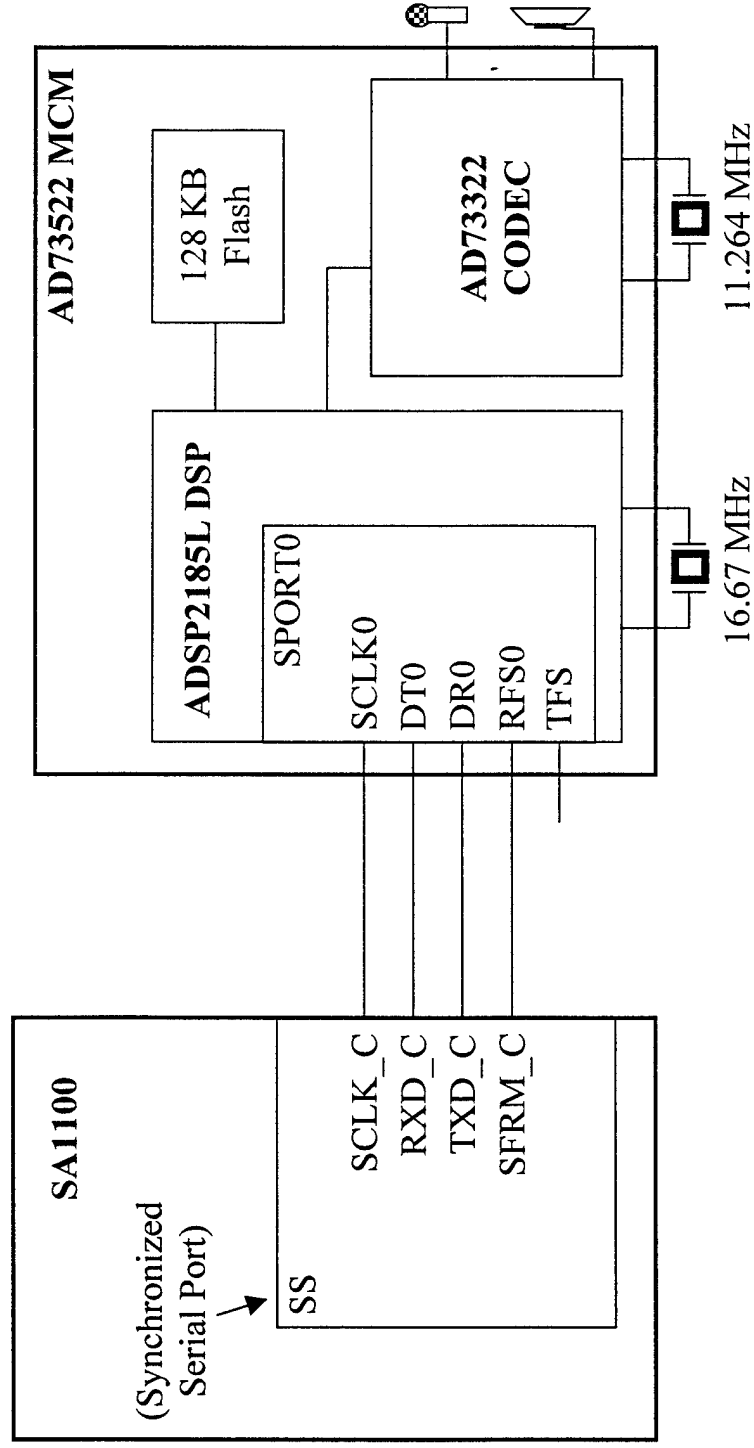


Fig. 7.10 Hook-Up of Audio Multi-Chip Module to SA1100

Voice Compression and Noise/Echo Cancellation

The tasks to be achieved on the WID to guarantee best voice-recognition results while keeping the voice bandwidth to a minimum are:

- noise cancellation
- echo cancellation
- voice compression

As discussed above, it appears that Configuration 4 in Fig. 7.2 may be the most advantageous and practical from the point of view of both the hardware and the software development. From the point of view of the hardware, it is advantageous to adopt a distributed voice processing approach. As shown in Fig. 7.11, therefore, the noise cancellation is performed by the microphone, the echo canceling software is ported to the DSP, and the voice compression is done by the SA1100. L&H's Automatic Speech Recognition (ASR) engine also performs some noise cancellation. In this way, no system component is overloaded.

From the point of view of the software, L&H has expressed interest in helping ViA port their software to the WID. In addition, L&H was already planning to port their software to the same processors used by the WID. Therefore, while ViA will keep monitoring alternatives, as shown in Table 7.3, the development effort over the next few months will concentrate on L&H's software.

Table 7.3 Sizing of DSP Memory Requirements				
	MIPS	Flash (kB)	Program RAM (kB)	Data RAM (kB)
Available on the AD73522	33	128	48	32
Option 1				
NCT Echo Canceling (on AD73522)	12	11.2	5.2	6
Voxware Voice Compression (on SA1100)	17	60	4	4
Net on AD73522:	4	56.8	38.8	22
Option 2				
L&H Echo Canceling (on AD73522)	28	5.2	6	0.2
L&H Voice Compression (on SA1100)	6	10.4	5	6.4
Net on AD73522:	5	122	42	31




Fig. 7.11 Distributed Voice Processing for WID + Wearable Computer

Chapter IX

PRELIMINARY PHYSICAL MODEL OF THE WID




Fig. 9.1 WID with Siliscape Display, Built with 1998 Components. Inches (mm)




Fig. 9.2 WID with Siliscape Display, Built with 1999 Components. Inches (mm)




Fig. 9.3 Rendition of Projected Appearance of WID Prototype